The Map and the World

Elements of a Philosophy of Cartography

Silvano Salvador

Preliminary Note

This volume stands apart from the two technical series it draws on — the cartographic distortion series (No Map Is Best) and The Impossibility Theorems of Cartography — while depending on both. Those volumes establish, with full mathematical rigor, a set of results about what representation of a curved surface on a flat one can and cannot achieve: Tissot's theorem on the incompatibility of conformality and equal-area representation, the topological obstruction expressed by the hairy ball theorem, the retroazimuthal problem and its historical solutions, the recent modernization of the U.S. geodetic datum. This volume does not re-derive any of those results. It takes them as established and asks a different kind of question of them: not whether they are true, but what it means, philosophically, that they are true — for representation in general, and for the specific human practice of finding one's place and one's direction.

The danger of this kind of volume is well known and is stated here rather than concealed: technical results, once removed from their proofs, are easily turned into loose metaphor. A theorem about vector fields on a sphere is not, by itself, a theorem about the limits of any 'global' human orientation, ethical or epistemic. Where this volume uses a result as a rigorous constraint, it says so and cites the chapter of the source volume where the proof is given. Where it uses a result only as an image or an analogy — a legitimate and old practice in philosophy, but a different one — it says that too, explicitly, rather than letting the two registers blur into each other. The reader is owed the distinction at every step, not only in a preface.

This is also, by intention, a shorter and more discursive volume than its technical companions. It does not aim at completeness in the philosophical literature it touches — representation, convention, phenomenology of space, the politics of maps each have vast literatures of their own — but at a single sustained argument, built honestly on the specific mathematical results this author has already established elsewhere, and nowhere else.

Chapter 1 — Introduction: Why Cartography Is a Philosophical Problem

1.1 A practical craft with a theoretical wound

Cartography presents itself, for most of its history, as a practical craft: a set of techniques for producing a usable flat picture of a curved world. Sailors, surveyors, and states wanted maps that worked, not maps that were philosophically innocent. And yet a wound sits at the center of the craft from the start, one that no amount of technical refinement can close, because it is not a technical limitation but a mathematical fact: a curved surface cannot be flattened without distortion. This is not a claim about the current state of cartographic technique. It is a claim, provable in the strict sense, about the relationship between two classes of object — surfaces of constant positive curvature and the Euclidean plane — that holds independently of who is doing the mapping or with what instruments.

The companion volumes to this one give the proof its full technical treatment: Gauss's Theorema Egregium establishes that curvature is an intrinsic invariant, unreachable by any bending that does not stretch or tear; Tissot's indicatrix quantifies, at every point of every possible projection, exactly how much and in what way the local picture must lie. What this volume asks is what follows from taking that proof seriously outside cartography — as a fact about representation as such, wherever a lower-dimensional or otherwise simplified picture is asked to stand in for a richer structure it cannot fully contain.

1.2 The two temptations

Two temptations present themselves immediately, and this volume tries to resist both.

The first is deflationary: to say that of course every map distorts, that this is trivial, that nothing philosophical follows from a fact so obvious it barely needs

restating. This response mistakes the ubiquity of a fact for its triviality. That every map distorts is indeed easy to state; what is not trivial is the precise structure of that distortion — that it obeys an exact trade-off law (Chapter 4), that certain configurations of the trade-off are topologically forbidden altogether rather than merely difficult (Chapter 5), that the choice among the surviving options is genuinely underdetermined by the geometry itself and requires an additional, non-geometric criterion (Chapter 6). None of this is available to intuition in advance of the proofs; all of it is philosophically substantive once proved.

The second temptation is inflationary: to treat the cartographic result as license for an unrestricted metaphor of distortion, applicable by loose analogy to any act of representation whatsoever — language, memory, ideology, the self. This is the move this volume's Preliminary Note already warned against. The mathematics constrains flat pictures of curved surfaces; it does not, by itself, constrain anything else. Where this volume extends the discussion beyond cartography strictly speaking, it will try to show a specific structural analogy — a shared reason why distortion is forced, not merely a shared word — rather than assert one by rhetorical momentum.

1.3 What this volume actually argues

The argument of this volume, stated in advance, runs as follows. A map is a special and unusually well-understood case of a more general situation: a structure-preserving correspondence is sought between a richer object and a poorer one, and the mathematics of that situation forces a choice among incompatible virtues rather than permitting all of them at once. Cartography is valuable to philosophy not as a metaphor but as the one case where this general situation has been worked out with complete rigor, so that the shape of the trade-off — what exactly is preserved, what exactly is sacrificed, and why no cleverness can avoid the sacrifice — is fully visible rather than merely asserted. Part II of this volume treats the trade-off itself, and the topological cases where even a partial solution is forbidden. Part III turns from the abstract trade-off to

the situated, first-person question of orientation — what it is to find one's direction toward a fixed point, literally in the case of the qibla and structurally in any case of orienting oneself toward a center one does not occupy. Part IV asks what follows from taking the inevitability of distortion as a starting condition of representation rather than as a defect to be regretted, and closes with the question of what, if anything, the cartographic case teaches other representational practices — asked cautiously, and only after the specific structural analogy, not the word alone, has been shown.

None of this requires the reader to take any result of the companion volumes on faith. Each technical claim used here is cited to the chapter where it is proved, and the proof is not repeated. What is added here is the question those volumes do not ask: given that this is how flat pictures of curved worlds must behave, what does it tell us about the wanting of a picture at all?

1.4 Why the Philosophical Use Requires the Mathematical One

The Preliminary Note already stated the discipline this volume tries to hold to: a rigorous constraint is cited to the chapter of the source volume that proves it, an image or analogy is marked as such, and the two registers are never allowed to blur. It is worth pausing, before the volume proceeds, on why this discipline matters as much as it does, because the alternative is not simply looser writing but a specific and common failure this volume calls epistemic borrowing: an author reaches for a genuinely proven result, uses its vocabulary and its air of settled fact to dress an intuition that has not itself been proven, and lets the reader's confidence in the theorem transfer, unearned, to the intuition riding along with it. The hairy ball theorem is proved; the claim that 'no ideology can orient a society everywhere consistently' is not, and stating both in the same breath, with the same tone of mathematical inevitability, invites the reader to grant the second the certainty that belongs only to the first. This is not a hypothetical risk. A great deal of the loosest writing about maps, some of it examined directly in Chapter 2's genealogy of 'the map is not the territory', trades on exactly this transfer, and no small part of what gives Baudrillard's

inversion of Borges its rhetorical force is the sense, never quite made explicit, that a genuine mathematical impossibility stands behind an argument that is in fact entirely about signs and their relation to a reality they no longer represent.

This volume's response is not to avoid the analogy, since Parts III and IV depend on it, explicitly, for their entire second half, but to make the transfer visible every time it happens, so that the reader can evaluate the analogy on its own philosophical merits rather than on the borrowed authority of a proof it does not actually extend to cover. The cost of this discipline is a certain repetitiveness: the reader will notice, chapter after chapter, the same distinction being redrawn between what a companion volume proves and what this volume is doing with it. That repetitiveness is not a stylistic failure to be smoothed away in a later revision. It is the argument's own honesty mechanism, and removing it would make the book easier to read and harder to trust.

1.5 A Worked Example: Two Sentences That Look Alike

Section 1.4 named epistemic borrowing in the abstract; it is worth showing, once, exactly how to tell a licensed use of a theorem from an unlicensed one, so that the distinction is a working tool for the rest of this volume and not only a promise made in its opening pages. Consider two sentences, each plausible on a first reading, each reaching for the same cartographic vocabulary. The first: 'no projection can be simultaneously conformal and equal-area, and a Mercator chart therefore misrepresents area at high latitudes by a computable, bounded factor.' The second: 'no representation can capture every aspect of what it represents, and a biography therefore misrepresents its subject by a comparable, if less computable, factor.' The two sentences share a grammatical shape and even a rhetorical cadence; only the first is a claim this volume can defend without qualification, because only the first names a quantity — the scale factors a and b of Section 4.1 — for which a theorem specifies an exact, provable relationship. The second sentence borrows the first's air of exactness ('a comparable... factor') for a claim that has no such quantity behind it: no theorem specifies what a biography must sacrifice, in what unit, in exchange for what, and 'comparable' is

doing exactly the illegitimate work the Preliminary Note warned against, smuggling a false precision into a domain where only a much vaguer and more defensible claim — that biographies, too, must select and omit — is actually available.

The test this contrast yields, and the one this volume will apply silently from here on wherever a comparison to cartography is made, is simple to state though not always simple to apply: ask whether the sentence, stripped of its cartographic borrowing, still names a specific quantity governed by a specific proof, or whether the proof has been left behind at the border and only its confidence has crossed over. Chapters 10 through 12 will apply this test explicitly and at length to the philosophy of scientific models, where the temptation to borrow is strongest because the vocabulary — model, representation, idealization — is shared almost word for word with cartography's own. Flagging the test here, before any of that material is reached, is meant to let the reader apply it independently, chapter by chapter, rather than trusting the volume's own application of it on faith.

1.6 The Wound Before the Proof: A Brief History of Suspecting What Gauss Would Prove

Section 1.1 named Gauss's Theorema Egregium as the proof that closes off, permanently, any hope of an isometric world map, and it is worth pausing, before Chapter 2 turns to the phrase 'the map is not the territory' and Chapter 4 states Gauss's result in full, on a fact this volume's framing has so far left implicit: the wound was felt, acutely and specifically, for well over two thousand years before it was proved. Marinus of Tyre and Ptolemy both wrestled, in the second century, with the specific problem of representing a spherical Earth's meridians and parallels on a flat chart without an accepted account of why no arrangement could ever get every distance right at once; Ptolemy's own conic construction, and Marinus's cruder equirectangular alternative, are documented attempts to minimize a distortion neither author had any way to prove was unavoidable in principle, because the differential-geometric vocabulary in which

such a proof could even be stated, curvature as an intrinsic invariant, did not exist and would not exist for another sixteen centuries.

The interval between suspicion and proof is itself philosophically instructive, and instructive in a way that bears directly on this chapter's opening claim that cartography's wound is 'not a technical limitation but a mathematical fact.' For most of that interval, the absence of an isometric world map was indistinguishable, from the working cartographer's vantage point, from an unsolved technical problem, the kind of thing better instruments or a cleverer construction might eventually fix; every named projection devised across those centuries, and there were dozens, can be read as an attempt, made in good faith, to be the exception. Gauss's 1827 Disquisitiones generales circa superficies curvas changed the status of every one of those attempts retroactively, without invalidating any single one of them as a projection: it did not show that Mercator, or Ptolemy's conic, or any predecessor had made a correctable error, but that the entire class of attempts had been aimed, unknowingly, at a target the curvature of the sphere makes strictly unreachable, for a reason having nothing to do with anyone's ingenuity and everything to do with an invariant no bending of the surface can alter.

This history sharpens rather than merely illustrates the distinction Section 1.1 draws between a craft frustration and a mathematical fact, because it shows the two states occupying the same historical space at different times without changing what was, all along, actually true. The sphere's Gaussian curvature was already nonzero, and already an obstruction to isometry, in Ptolemy's century as much as in Gauss's; what changed in 1827 was not the geometry but the available proof of it, and with the proof, the epistemic status of every future attempt at an isometric projection, from a live technical hope to a settled impossibility. The lesson this volume draws from the interval is not that suspicion before proof was irrational — the persistent, centuries-long failure to find an isometric projection was itself, statistically if not deductively, decent evidence something structural was blocking every attempt — but that suspicion and proof are different achievements, and that this volume's own discipline, distinguishing what a

companion volume has actually proved from what remains, at any given chapter, offered only as a plausible reading, is the same discipline the history of cartography spent two thousand years learning to need.

1.7 Kant's Transcendental Schematism: The Image and the Rule

Kant's first Critique distinguishes an image (Bild) from a schema: a schema is not a picture of a triangle but a rule for constructing any triangle, mediating between a bare concept and the particular intuitions that can fall under it. The opening chapters of this volume have already needed something like this distinction without naming it — Section 1.1 treated a map neither as a bare particular picture of one place nor as an abstract geometric rule floating free of any territory, but as something that stands between the two, a procedure that meets sensibility partway.

Kant's worked example is arithmetic: the schema of the number five is not any one image of five dots but the rule for successively adding units that could generate any such image. A cartographic projection is, in the strict and not merely suggestive sense, exactly this kind of rule: Section 3.2's distinction between a map, a model, and an isomorphism already treated a projection as a function, and a function is a Kantian schema in miniature — not a sheet of paper but the systematic procedure that generates any particular sheet from any particular territory, the general rule of which Chapter 4's named projections (Section 4.4) are individual instances, images in Kant's sense.

The analogy has a clear stopping point, and the Preliminary Note's discipline requires stating it. Kant's schematism is transcendental: it purports to explain how any concept whatsoever can meet any possible intuition, a condition of experience in general. A cartographic projection formula explains how one narrow class of concepts — positions on a curved surface — can meet one narrow class of intuitions — marks on a flat sheet. What survives the comparison is not Kant's full transcendental machinery but a single structural point worth keeping: the wound described in Section 1.1 belongs to the rule, not to any of its

images, and no amount of care in drawing a particular sheet can repair a defect that sits one level up, in the schema itself.

1.8 Ptolemy's Grid: The Wound's First Systematic Recognition

Section 1.6 sketched a brief history of suspecting, before any proof existed, that flattening a curved world exacts a cost. It is worth pausing on the single earliest systematic response to that suspicion: Claudius Ptolemy's Geographia, which does not merely list places by latitude and longitude but devotes an opening methodological book to comparing rival projections — a simple conic and a more elaborate pseudoconic construction — explicitly on the grounds that neither preserves every relation a reader might want, and that the choice between them is a choice about which relations to sacrifice.

This matters for Section 1.6's history because Ptolemy is not merely aware that maps distort, which Section 1.2 already treated as too obvious to be philosophically interesting on its own; he is aware, without any of the machinery Chapter 4 would eventually supply, that the distortion has structure — that different constructions distort differently, trading one kind of fidelity for another rather than simply being better or worse versions of the same thing. His second projection is offered explicitly as an improvement in appearance at some further, named cost in ease of construction, a trade-off judgment in the same family as Section 4.7's discussion of Robinson's much later compromise projection, arrived at by inspection rather than proof.

What Ptolemy lacks, and what distinguishes his moment from the one Chapter 4 describes, is any way to know whether he has found the best available trade-off or merely a serviceable one, or whether some further projection might dominate both of his. That gap — between recognizing that a trade-off exists and being able to characterize its exact shape — is precisely the gap this volume's companion series closes, and naming Ptolemy's achievement carefully, as the former without the latter, keeps Section 1.6's history honest about what had and had not yet been shown.

1.9 A Second, Independent Wound: The Coastline Paradox

Sections 1.1 through 1.8 have treated a single wound, the one Gauss's Theorema Egregium proves: curvature forbids an isometric flattening of the sphere. It is worth noting, before Part II develops that wound in full technical detail, that it is not the only mathematical difficulty a map faces, and that a second, genuinely independent one was identified by the meteorologist Lewis Fry Richardson in the 1950s and given its modern name and explanation by Benoit Mandelbrot's 1967 paper 'How Long Is the Coast of Britain? Statistical Self-Similarity and Fractional Dimension.' Richardson, attempting to correlate the length of national borders with the frequency of war between neighboring states, found that the measured length of a coastline or a land border is not a fixed number at all: measuring Britain's coast with a longer ruler yields a shorter total length than measuring it with a shorter one, and the measured length grows without evident bound as the ruler shrinks, because a coastline's irregularity repeats, self-similarly, at every scale a ruler could be chosen to probe.

Mandelbrot's contribution was to show this was not measurement error but a genuine mathematical property of the curve itself, formalized as a fractal dimension strictly between one and two — a coastline is, in this precise sense, too irregular to be a one-dimensional curve and too sparse to fill a two-dimensional area, and its 'length' is consequently not a well-defined number independent of the scale at which it is asked for, the way the circumference of a circle is. This is a wound of a different kind from Section 1.1's: Gauss's obstruction concerns curvature, a property of a smooth surface, and forces distortion in the relation between a curved surface and a flat one regardless of the map's scale; Richardson and Mandelbrot's obstruction concerns self-similar irregularity, a property some real coastlines happen to have at every scale, and forces indeterminacy in a single quantity, length, that has nothing to do with the projection chosen and would afflict even a perfect, distortion-free scale model.

The two wounds are worth keeping exactly as separate as the Preliminary Note asks every technical claim in this volume to be kept from its neighbors. Part II's

theorems, Tissot's and the hairy ball's alike, are true of every map regardless of what is being mapped, because they follow from the sphere's geometry and topology alone; the coastline paradox is true only of the specific, empirically self-similar geographic features it has been measured for, and a sufficiently smooth or straight border, a survey line running along a meridian, raises no such difficulty at all. What the two share, and the only claim this section makes on their behalf together, is the general lesson Section 1.2 already stated in the abstract: a representation's failure to capture its territory fully is not one uniform fact awaiting one uniform diagnosis, but a family of distinct, independently provable obstructions, each with its own precise scope, and this volume's discipline requires knowing, at every point, which member of the family is actually in play.

1.10 Scale as a Third Axis: Generalization and the Cartographer's Rule of Thumb

Sections 1.1 through 1.9 have identified two independent wounds, curvature and self-similar irregularity, that no choice of projection or instrument can heal. A third constraint belongs beside them, not because it is a theorem in the strict sense the other two are, but because working cartographers have treated it as a near-inviolable rule of practice for as long as printed maps have existed, and because it interacts with both of the proven wounds in a way worth making explicit before Part II turns to the first of them in full technical detail. The constraint is generalization: the fact that a map drawn at one scale cannot simply be the same map redrawn smaller at another scale, because features that read as informative detail at one scale become illegible clutter at another, and a coastline whose length is already scale-dependent, per Section 1.9, must additionally be simplified, smoothed, and selectively omitted as the map's scale shrinks, not merely shrunk uniformly along with everything else.

The practical rule cartographers have used to manage this, sometimes called Tobler's radical law after the geographer Waldo Tobler's proposed formula relating the number of features retained to the square root of the scale ratio, is

not a theorem of the kind Sections 1.1 through 1.9 have been careful to distinguish from opinion; it is closer to a rule of thumb, empirically calibrated against what legible maps have actually done, than to a derivation from first principles. But its existence is philosophically instructive precisely because it shows the wound propagating into a third domain. Curvature forces distortion in angle, area, or both, at every scale simultaneously; self-similar irregularity makes a coastline's own length scale-dependent; generalization adds a further, editorial layer of scale-dependent loss that has nothing to do with either — a perfectly conformal, perfectly measured map of a coastline, drawn at continental scale, still cannot show every inlet the same map would show at the scale of a single harbor, not because the geometry forbids it but because the paper, the eye, and the map's purpose cannot absorb it.

What generalization adds to this chapter's argument, then, is not a fourth theorem to set beside Gauss's and Mandelbrot's, but a reminder that the wound the chapter has been describing is compounded, in ordinary cartographic practice, by constraints that are real, load-bearing, and irreducible to any single mathematical result, without being provable in the sense the rest of this volume insists on before it will call something a limit rather than a habit. The discipline this volume has committed itself to, of naming exactly which obstruction is in play at each point, requires saying plainly that this one belongs to a different, humbler category than curvature or fractal dimension: a constraint of design and legibility, not of geometry, but no less real to the cartographer holding the pen for that.

1.11 A First Encounter with the Wound in Ordinary Life: Why Airline Route Maps Curve

Before the volume turns to Gauss's theorem in Chapter 4, or to any of the historical episodes of Sections 1.6 through 1.9, most readers have already met the wound this chapter describes, without necessarily recognizing it as such, in an ordinary and mildly puzzling piece of consumer technology: the route map shown on an airline's in-flight display or ticketing website for a long transoceanic

flight, on which the drawn path from, say, New York to Tokyo bows northward in a broad arc toward the Arctic rather than running the straight westward line a flat mental picture of the globe would suggest. The arc is not an error, an evasion, or a marketing flourish; it is the visible trace of exactly the wound Section 1.1 named, projected outward from the abstract geometry of a curved surface into a document millions of travelers glance at without further thought. A great-circle route, the shortest path between two points on a sphere, is a straight line only on the sphere itself; represented on any flat map, it must bend, and the direction and severity of the bending depend entirely on which of the many admissible flattenings, Chapter 4 will formalize the choice as a trade-off rather than a menu of free options, the mapmaker happened to choose for that particular chart.

What makes the airline route map a genuinely useful first encounter with the book's central wound, rather than merely a curiosity, is that it demonstrates a claim Section 1.3 will state more carefully: the distortion is not a defect correctable by a better cartographer or a higher-resolution map, but a structural consequence of representing a curved surface on a flat one, present in some form in absolutely every flat map ever drawn, including the one on the seatback screen. A traveler who has ever wondered why a flight to Asia crosses so far north, or why a great-circle route looks so counterintuitive on a familiar rectangular world map, has already, without any of the mathematics this volume will supply starting in Chapter 4, encountered the wound Section 1.1 opened this chapter by naming; the rest of this volume is, in one sense, no more than an extended and disciplined answer to the puzzlement that route map already provokes.

Chapter 2 — "The Map Is Not the Territory": A Critical Genealogy of a Commonplace

2.1 A phrase that outran its argument

Few phrases have traveled farther from their original context than 'the map is not the territory.' Coined by Alfred Korzybski in the 1930s as part of a broader program (general semantics) aimed at correcting habits of language that treat words and their referents as interchangeable, the phrase has since become a free-floating commonplace, invoked to license almost any claim that a representation falls short of what it represents. This chapter does not attempt a full history of the phrase's reception; it asks a narrower question directly relevant to this volume — what, if anything, the phrase means once it is held to the standard the cartographic result actually provides, rather than the standard of a well-turned slogan.

Korzybski's own point, stripped of the slogan's later drift, is a claim about confusing levels of abstraction: a word is not the thing it names, a model is not the modeled system, and treating the two as identical produces specific, diagnosable errors of reasoning. This is a claim about a category mistake, not, in Korzybski's original formulation, a claim about geometry. The cartographic case this volume works from is different in kind: it does not say merely that a map is not identical to its territory (a claim true of every representation whatsoever, trivially, and not distinctive of maps) but that a map cannot be made to preserve every property of its territory simultaneously, for a provable geometric reason specific to the curvature of the sphere. The slogan is compatible with a world in which, given enough cleverness, the gap between map and territory could in principle be closed. The theorem is not.

2.2 Borges's fable and its inversion

The most famous literary treatment of the theme, a very short prose fragment attributed by Jorge Luis Borges to a fictional seventeenth-century source, imagines a map drawn at the same scale as the empire it depicts, so that the map and the territory coincide point for point — and then imagines this map abandoned as useless, its tattered remains later found sheltering animals in the desert. The fable is usually read, correctly, as a satire of an impossible and pointless ideal: perfect fidelity purchased at the price of all utility. What is worth

adding, from the standpoint of this volume, is a further and stricter sense in which the one-to-one map is impossible, not merely useless. Even a map drawn at the territory's own scale, laid flat, would still have to distort — because the territory in question is (for any cartographic case of genuine interest) a piece of a curved surface, and no flat laying-out of a curved surface preserves all of its metric properties, regardless of scale. Borges's fable stages the pointlessness of perfect fidelity; the cartographic theorems stage its impossibility. The two are not the same criticism, and the second is the sharper one, because it survives every objection that grants the map-maker unlimited resources, unlimited scale, and unlimited care.

2.3 Baudrillard's inversion and where this volume parts from it

Jean Baudrillard took up Borges's fable as an epigraph for an argument that runs, in effect, in the opposite direction from the one this volume pursues: that in a media-saturated condition the relation reverses, so that models and signs precede and generate what is taken for the real, and the territory itself decays and disappears while the map endures — the map is no longer poorer than the territory but comes to substitute for it. This is a genuinely different problem from the one this volume treats. Baudrillard's concern is with the social and semiotic status of representation under specific historical conditions; it is not a claim about the geometry of representation and does not depend on, or bear on, whether a projection can be simultaneously conformal and equal-area. Citing him here is meant to mark a boundary, not to borrow authority: this volume's problem is prior to and narrower than Baudrillard's, and does not inherit his conclusions merely because both discussions use the word 'map.'

2.4 What survives the genealogy

What should survive, for the purposes of this volume, once the commonplace is traced back through Korzybski, Borges, and Baudrillard, is not the slogan itself but a sharpened question it gestures toward without answering: under exactly

what conditions, and to what exact degree, must a representation fall short of what it represents, and is that shortfall contingent on the representer's skill or forced by the structure of the objects involved? Korzybski's version answers with a warning about confusing categories. Borges's version answers with a fable about futility. Neither offers a proof. The cartographic case this volume works from offers exactly that: an exact, quantified, provable answer to the question of how much a flat picture of a curved surface must lie, and about what, specifically — the subject of Chapter 4.

2.5 What Korzybski Actually Claimed, and What the Slogan Dropped

The genealogy this chapter has traced so far moves from Korzybski's coinage to Borges's fable to Baudrillard's inversion, treating 'the map is not the territory' largely as a free-floating phrase whose subsequent life outran its origin. It is worth returning briefly to that origin, because Korzybski did not intend the phrase as a piece of loose common sense about representation in general; it was a technical claim embedded in a larger and considerably more demanding apparatus, general semantics, built around what he called the structural differential and a theory of levels of abstraction running from an unspeakable physical event, through a nervous system's non-verbal abstraction of it, up through successive verbal descriptions each abstracting further from the last. Within that apparatus, 'the map is not the territory' was not primarily a warning against confusing representation with reality in the abstract; it was a specific diagnostic claim that a great deal of human confusion, and in Korzybski's own framing psychological distress, arises from confusing one level of abstraction with another, mistaking a description of a description for a description of the thing itself, or reacting to a word as though it carried the properties of what it names.

Almost none of this survives in the phrase's later, popular life, and the loss is not incidental to the genealogy this chapter is tracing. What survived was the four words and their negative form, stripped of the structural-differential apparatus that gave the original claim its content and its testable structure. What this leaves

behind is a slogan that sounds like an insight but functions, in most of its later uses, as an assertion with no proposition attached: of course a representation is not the thing represented, and stating this fact in isolation, without Korzybski's account of levels and where confusion between them actually occurs, explains nothing about when or why the distinction matters in a given case. The sections that follow, on Borges and on Baudrillard, can now be read with this specific loss in view: both authors inherit the slogan's rhetorical force while discarding, along with everyone else, the diagnostic machinery that once made it more than an aphorism.

2.6 The Phrase in Managerial and Therapeutic Discourse: The Commonplace at Its Furthest Remove

What Section 2.4 called the commonplace's 'free-floating' afterlife has, in the decades since Korzybski, drifted further still, into a register neither philosophical nor literary but managerial and therapeutic, where 'the map is not the territory' now functions as a stock opening line in leadership seminars, coaching manuals, and the neuro-linguistic programming literature that adopted it as one of its founding presuppositions. It is worth pausing on this final drift, briefly, not because it deserves the sustained treatment given to Borges or Baudrillard but because it is the clearest available specimen of exactly the failure Section 2.1 diagnosed in the abstract: a phrase surviving its argument entirely, retained for its cadence long after the structural claim that once gave it content has been dropped. In this literature the slogan typically supports a claim of the form 'your beliefs about a situation are not the situation itself, and changing the map can change your experience of the territory' — a claim about the malleability of perception and belief that Korzybski's own apparatus, built around a fixed hierarchy of levels of abstraction running from an unspeakable physical event up through increasingly general verbal labels, was in fact designed to discipline, not license: general semantics warned against treating a low-level label as though it exhausted the event it named, which is nearly the reverse of counseling a reader to feel free to swap labels because none of them touch the event anyway.

This volume's interest in the specimen is narrow and instrumental rather than critical of the coaching literature on its own terms, which is not this volume's business to adjudicate. What the managerial drift shows, with unusual clarity because the distance traveled is so large, is that a phrase's rhetorical usefulness is almost entirely uncorrelated with how much argument still travels with it. Korzybski's phrase began attached to a specific, if contestable, theory of levels of abstraction; it now most commonly appears attached to no theory at all, functioning instead as a permission slip for whatever claim about the flexibility of perspective a given speaker wishes to make next. The cartographic case this volume works from cannot suffer the same fate by the same route, because Chapter 4's theorem is not a phrase that can be repeated without its proof; a reader can forget why Tissot's theorem is true and still be prevented, by the theorem itself, from misstating what it says, in a way no amount of repetition protects a slogan from.

2.7 Bateson’s Frame: The Map-Territory Distinction Turned Into a Theory of Mind

Section 2.5 recovered what Korzybski actually claimed, against the slogan’s later flattening, and Section 2.6 traced that flattening into managerial and therapeutic cliche. Between Korzybski’s own general semantics and the phrase’s eventual dilution sits a third episode this chapter’s genealogy has not yet named: Gregory Bateson, who studied under Korzybski’s influence and carried the map-territory distinction into cybernetics and, from there, into a general theory of mind. Bateson’s own restatement, in Steps to an Ecology of Mind, sharpens rather than loosens Korzybski’s original point: ‘the map is not the territory, and the name is not the thing named,’ but Bateson’s interest is specifically in what a map preserves, not merely in reminding the reader that something is lost. His famous formulation of information as ‘a difference that makes a difference’ treats a map, or any representation, as valuable exactly to the extent that the differences it encodes correspond to differences in the territory that matter to some receiving system — a criterion of selection, not merely of loss.

This is a genuine extension of Section 2.4's 'what survives the genealogy,' not a repetition of it, because Bateson supplies a mechanism the slogan by itself does not: his account of Logical Types, borrowed from Russell and Whitehead's theory of types in logic, treats the map-territory relation as one instance of a general hierarchy in which a class cannot be a member of itself, a name cannot be the thing named, and, crucially, a message about a message (Bateson's central case is play: a nip that signifies a bite, and the further frame, itself communicated, that the nip is not to be taken as the bite it resembles) occupies a logically distinct level from the message itself. Applied back to cartography, this supplies a way of stating why the confusion the slogan warns against is not merely careless but a specific logical-type error: treating a symbol at one level (the map's blue line) as if it possessed properties that belong only at the level it represents (the river's wetness, its capacity to drown a swimmer) conflates levels Bateson's hierarchy keeps formally separate, in a way close to, but more general than, Section 3.4's distinction between a map's symbol and its target.

Bateson's frame is worth holding apart, in the end, from Baudrillard's inversion in Section 2.3 as sharply as Section 2.4 held Korzybski apart from the slogan's later users, because the two post-Korzybski trajectories run in opposite directions. Baudrillard's simulacrum dissolves the map-territory distinction from the side of the map, arguing the represented eventually disappears beneath its own representations; Bateson's logical typing preserves and formalizes the distinction from the side of information, treating the difference between levels not as a wound to be lamented or a hierarchy to be inverted but as the very structure that allows a map, a play-nip, or a scientific model to communicate anything at all. Cartography's own theorems, proved without reference to either author, side unambiguously with Bateson's reading: Tissot's theorem states precisely which differences a projection preserves and which it cannot, a logical-type-respecting fact about levels, not a symptom of the representation swallowing its territory whole.

2.8 Deleuze and Guattari's Rhizome: The Map That Is Not a Tracing

Sections 2.2 through 2.8 have traced a genealogy in which the map, from Borges's fable onward, is repeatedly cast as the party under suspicion: too faithful in Borges, too autonomous in Baudrillard, too readily mistaken for the territory in Korzybski's original warning and its managerial afterlife. Gilles Deleuze and Félix Guattari, in the introductory plateau of Mille Plateaux, invert this posture so completely that the inversion is worth setting down as its own datum in the genealogy rather than folding it into the critiques already surveyed. For Deleuze and Guattari, the map is not the suspect term at all; it is the tracing, decalque, that stands accused, because a tracing merely reproduces a structure already given in advance, projecting a pre-existing order back onto whatever it is applied to, while a map — in their technical, deliberately extended sense of the word — is open, experimental, without a fixed center, oriented toward connection and performance rather than toward the reproduction of an origin.

The distinction does not restore any of the properties Chapters 4 and 5 proved cartography's actual maps cannot have; Deleuze and Guattari are explicit that their map is not the paper artifact this volume has been examining, but a figure for a mode of thought, rhizomatic rather than arboreal, and the extension from paper chart to philosophical figure is precisely the kind of move Section 2.4 already flagged as needing separate justification each time it is made. What the rhizome image does contribute, against the rest of this chapter's genealogy, is a demonstration that the map/territory or map/tracing pairing is not intrinsically loaded in one direction; Korzybski's warning treats the map as the thing likely to be mistaken for more than it is, Baudrillard's inversion treats it as the thing that has replaced what it was meant to represent, and Deleuze and Guattari treat it as the more open, less presumptuous of the two terms, with the tracing cast as the one that mistakes a single, already-fixed structure for the whole of what could be represented.

Read against this volume's own argument, the rhizome passage is best treated as a caution rather than a resource: it shows how readily the map/territory vocabulary can be redeployed to support opposite evaluative conclusions once it leaves the domain, actual paper cartography under Tissot's and the hairy ball's

constraints, where its terms have a fixed technical meaning. Section 2.4's question, what survives the genealogy, gains a further answer here: what survives is not any particular verdict on whether maps are dangerously faithful or admirably open, since the vocabulary alone will underwrite either verdict on demand, but the narrower, load-bearing fact this volume has kept returning to beneath all of them — that an actual map is a homomorphism under a real geometric constraint, and every philosophical extension of the phrase, celebratory or cautionary, is answerable to that fact rather than a substitute for examining it.

2.9 Borges's Map Taken Literally: What a One-to-One Correspondence Would Actually Cost

Section 2.2 read Borges's one-paragraph fable of the empire whose cartographers produce a map coextensive with the territory itself as a piece of rhetoric, a deliberately absurd image meant to needle the pretension of excessive fidelity. This section asks a more literal question, one this volume's own apparatus is actually equipped to answer: bracketing the fable's satire, could a one-to-one map exist at all, as a coherent representational object rather than merely an impossible-sounding image?

The answer follows directly from Section 3.2's vocabulary. A map is useful exactly to the extent that it is a homomorphism onto something poorer than the territory — fewer points, less detail, a lower dimension of information per unit distance — preserving some structure while discarding the rest. A same-scale copy laid over the territory it depicts preserves everything and discards nothing; it is not a homomorphism onto a poorer structure but the identity map onto the same structure, which is to say it has stopped being a representation of the territory and become, trivially, a second territory. Borges's empire did not build an unusually good map; it built a duplicate empire that happened to be made of paper, a distinction the fable's famous final image — the map left to rot in the desert, useless — already registers without quite stating it in these terms.

This matters for what survives the genealogy (Section 2.4). The fable is not, on this reading, a warning against Chapter 4's distortion in the geometric sense — scale is chosen prior to and independently of the Tissot trade-off, which concerns what happens once a scale smaller than one has already been fixed. Borges's real target is usability and finitude, not fidelity: a map is valuable because it is smaller and simpler than what it depicts, so that the demand for zero distortion, pushed to its logical limit, does not produce a perfect map but the abolition of mapping as such. The fable's lesson, taken literally rather than as ornament, turns out to agree with rather than anticipate Chapter 4's finding: some cost is not a regrettable feature of representation but the condition of there being a representation to speak of.

2.10 Eco's Sequel to Borges: Engineering the Impossible Map

Section 2.9 took Borges's fable of the one-to-one map literally and asked what a coherent version of the object would actually cost, concluding that scale one collapses representation into duplication. Umberto Eco, in a short essay titled (in its published English translation) 'On the Impossibility of Drawing a Map of the Empire on a Scale of 1 to 1,' had already asked a closely related but distinct question: not whether such a map is representationally coherent, but what would go wrong, physically and logistically, if an empire actually tried to build one.

Eco's essay works through several increasingly desperate engineering fixes — stretching the map over the territory it depicts, folding it, building it as a self-mapping surface — and shows each one collapsing into absurdity for reasons that are practical rather than geometric: the map-building material itself needs space, so the map competes with the territory for the room to exist, and any attempt to make the map thinner or more efficient only reproduces, at one remove, the original problem of representing a large thing compactly, the very problem a map is supposed to solve in the first place.

Read alongside Section 2.9, Eco supplies the practical half of an argument this volume has now made in full on both registers. Section 2.9 showed that a

one-to-one map fails as representation, becoming identity rather than depiction; Eco shows that even setting representation aside and asking only whether the physical object could be built and used, the attempt fails for an independent, entirely mundane reason — there is nowhere to put it. The fable's two most careful modern readings agree, from different directions, that the joke was never really about excessive fidelity at all, but about what 'map' stops meaning once compression is removed from its definition.

2.11 Debord's Dérive: The Map Refused Rather Than Reinterpreted

The genealogy this chapter has traced runs from a fable about a map too literal to be useful, through Baudrillard's inversion of that fable into a claim about a territory that has itself dissolved into simulation, to Bateson's and the therapeutic tradition's domestication of the phrase into a piece of self-help advice about not confusing one's beliefs with reality. There is a further station on this route that belongs here for a different reason: not because it extends the slogan's content, but because it stages a refusal of the map altogether, and the refusal is itself philosophically legible against everything Sections 2.1 through 2.10 have established. The Situationist practice of the dérive, associated above all with Guy Debord's 1956 essay 'Théorie de la dérive,' proposed an unplanned, drifting mode of urban movement explicitly opposed to the map's governing logic: where a map presupposes a stable, surveyable city that can be represented from above and navigated by reference to that representation, the dérive asked its practitioners to let the city's actual psychological and affective terrain, its 'psychogeography,' redirect their movement in ways no map could anticipate or record.

What makes this worth distinguishing from Baudrillard's move, treated in Section 2.3, is that Debord was not claiming the map had replaced or exceeded the territory; he was refusing the map's authority over lived movement through the territory, a different and in some ways more modest claim. The dérive does not assert that Paris has become indistinguishable from its representation, in the manner of Baudrillard's simulacrum; it asserts that a certain kind of

representation, the surveyed, gridded municipal map, systematically occludes a layer of urban experience, the psychogeographic, that a different practice of attention can recover. This is closer to Section 2.2's reading of Borges, in which the map's inadequacy is a matter of scale and purpose rather than metaphysical collapse, except that Debord's response is not Borges's cartographers' response of building a bigger map; it is to abandon cartographic navigation as a mode of urban self-location altogether, at least for the duration of the dérive, in favor of a form of attention the map was never designed to support.

The Situationist practice therefore occupies a position this chapter's genealogy has not yet named: a use of 'the map is not the territory,' or its practical equivalent, that neither corrects the phrase back toward Korzybski's engineering caution (Section 2.5) nor inflates it into Baudrillard's ontological claim (Section 2.3), but turns it into a program of deliberate disorientation, on the wager that what a map cannot show is sometimes worth experiencing precisely because it is unmapped. Whether this wager succeeds as urban practice is not a question this volume is positioned to adjudicate; what belongs to the genealogy is only the observation that the commonplace, once loose in the culture, generated not only theoretical elaborations and therapeutic domestications but at least one concrete practice built on refusing the map's premise rather than reinterpreting its content — a data point Section 2.6's survey of the phrase's furthest reaches did not include, and one that closes this chapter's genealogy on a note of practice rather than theory, appropriately, since Chapter 3 returns to the theoretical register the dérive stepped outside of.

Chapter 3 — Representation, Model, Isomorphism: The Map as an Epistemological Test Case

3.1 Why the map is a good test case and not just an example

Philosophy of science has a long-standing interest in the general question of what a model or a representation is, and how it relates to the system it

represents. That literature typically works with cases — scientific theories, diagrams, scale models, mathematical structures standing in for physical ones — chosen because they are important, not because they are simple. A cell diagram, a supply-and-demand curve, a wavefunction: each raises the general question of representation, but each also carries its own domain-specific complications, so that a claim about what representation requires is hard to separate from a claim about biology, or economics, or quantum mechanics in particular. The map is unusual among candidate test cases in that the relevant structure is exhausted by pure geometry. A projection is a function from a sphere (or a chosen approximation of the Earth's true shape) to a plane; what it preserves and what it distorts follow entirely from that function's mathematical properties. There is no residual empirical uncertainty about the target system, as there is in biology or economics, that could be blamed for a representation's shortfall. Whatever the map fails to preserve, it fails to preserve for reasons intrinsic to the relation between the two geometric objects — which makes cartography an unusually clean case for isolating what representation, as such, can and cannot do.

3.2 Three notions worth keeping apart

Discussions of representation tend to move between three related but distinct notions, and the cartographic case gives a reason to keep them apart rather than let them slide into one another.

The first is correspondence: a map assigns to every point of the sphere (or a suitable domain of it) exactly one point of the plane, and the assignment is well-defined and, for the projections of cartographic interest, continuous and invertible on its domain. Correspondence alone is a weak requirement — it is satisfied by an enormous space of possible functions — and by itself guarantees nothing about which properties of the sphere survive the passage to the plane.

The second is isomorphism in the strict sense: a correspondence that preserves a specified structure exactly. A map that is an isometry would preserve all distances; one that is conformal preserves all angles; one that is equal-area

preserves all areas. The cartographic theorems (Chapter 4) show that no projection of a non-flat region can be an isometry, and that no projection can be simultaneously conformal and equal-area. Isomorphism, in other words, is not merely difficult to achieve for the sphere-to-plane case: full isomorphism (isometry) is flatly impossible, and even the partial isomorphisms available are mutually exclusive.

The third is fidelity in the loose, non-technical sense typically invoked outside mathematics — a map is 'good' or 'accurate' if it serves its user's purposes, without any single preserved structure being specified. This is the notion actually at work when a subway map, wildly non-isometric and non-conformal, is nonetheless called an excellent map: it preserves topological connectivity and route order, which is what its users need, at the cost of every metric property, which they do not. Fidelity in this sense is not a weaker version of isomorphism; it is a different kind of success criterion altogether, indexed to a purpose rather than to a structure.

Confusing the third notion for the second is the most common error in casual talk about maps 'lying' or being 'accurate,' and it is worth naming as an error rather than a mere imprecision: a map can be maximally successful by the third standard while being, of mathematical necessity, a failure by the second. The philosophical interest of the cartographic case is that it forces this distinction into the open, because the second notion is provably and permanently unreachable in full, which means every actual map's success must be cashed out in terms of the first or the third.

3.3 What this suggests for representation beyond cartography

The suggestion this volume wants to defend, cautiously and only as a suggestion at this stage, is that the same three-way distinction — correspondence, strict structural isomorphism, purpose-relative fidelity — is doing unacknowledged work in other domains where 'model' or 'representation' is used loosely across all three senses at once. A scientific model of a physical system is very rarely

offered, or even intended, as a strict isomorphism; it is usually a correspondence that captures some specified structure (the dynamics relevant to a given question) while admitting, and sometimes advertising, its failure to preserve others. What the cartographic case adds, that a looser example cannot, is a fully worked demonstration — not an assertion — that in at least one important case, strict isomorphism is not merely hard to achieve but excluded on pain of contradiction, and that the resulting trade-offs among partial isomorphisms are themselves exactly classifiable rather than a vague residue of unmodeled complexity. Whether an analogous exact classification exists for other representational domains is a substantive question this volume does not attempt to settle; it can only report, honestly, that cartography settles it for the one domain in which it has actually been proved, and use that as a disciplined check on how much can be claimed elsewhere.

3.4 Homomorphism, Not Isomorphism: A More Exact Vocabulary

Section 3.2 distinguished representation, model, and isomorphism as three notions worth keeping apart, and used 'isomorphism' provisionally as a stand-in for the general idea of structure-preservation. It is worth being more exact, because the model-theoretic vocabulary already exists to say precisely what a map is and is not, and using it precisely removes an ambiguity that has quietly done philosophical damage elsewhere in the literature this volume engages. An isomorphism, in the strict sense, is a structure-preserving bijection: it maps every element of the source to a distinct element of the target, and back again, without loss. No map of a curved surface onto a flat one is an isomorphism in this sense, and Part II will shortly explain why none can be, but the more immediate point, available already from this chapter's own resources, is that a map was never a candidate for isomorphism to begin with, independent of any curvature theorem. A map necessarily discards information the territory contains: elevation reduced to a contour line's abstraction, or omitted entirely; the territory's history, its smell, the texture of its ground, absent by the nature of the medium before any question of projection arises. What a good map achieves, at best, is a homomorphism, a structure-preserving mapping that need not be one-to-one and need not be reversible, preserving some relations of interest, adjacency, relative bearing, sometimes relative area, while collapsing or omitting others entirely.

This distinction matters because 'the map is not the territory', read through it, stops being a truism and becomes a substantive question with a precise form: which relations does a given map homomorphically preserve, and which does it collapse or drop? That question has a different answer for a subway diagram, which preserves connectivity and sequence while discarding true distance and bearing entirely, and for a nautical chart, which preserves bearing and distance with unusual fidelity while discarding almost everything about what lies inland. Chapter 4's theorem about projections will show that even the more modest ambition of preserving angle and area together, let alone full isomorphism, is unavailable on a curved surface, but that theorem is best understood, with this

section's vocabulary in hand, as a further restriction on an already homomorphic, already lossy relationship, not as the sole source of a map's incompleteness.

3.5 Structural Realism and the Semantic View: Suárez, Giere, van Fraassen

Section 3.3 flagged, cautiously, a suggestion it did not attempt to settle: that the three-way distinction among correspondence, strict isomorphism, and purpose-relative fidelity might be doing unacknowledged work in the philosophy of scientific models. This section and the next enter that literature directly rather than merely gesturing at it, because a suggestion left unchecked is exactly the kind of move the Preliminary Note warned against. The semantic view of theories, associated with Suppes and developed by van Fraassen, treats a scientific theory not as a set of sentences but as a family of models, and asks a correspondingly narrower question of them than full isomorphism: a model is empirically adequate if some substructure of it, the part answering to what is actually observed, is isomorphic to the appearances, while the rest of the model is permitted to go proxy for the unobservable without any isomorphism claim being made on its behalf at all. This is already a retreat from the second notion of Section 3.2 toward something nearer the third, and it is worth naming as such rather than letting the word 'model' carry the analogy on its own.

Giere's perspectival realism pushes the retreat further, replacing isomorphism outright with similarity: a model represents its target to the extent that the two are similar in specified respects and to specified degrees, a relation Giere deliberately leaves graded and interest-relative rather than all-or-nothing. This is close, in spirit, to the homomorphism this volume defended in Section 3.4 over a stricter isomorphism — both give up the demand for a lossless correspondence — but similarity, unlike homomorphism, is not itself a fixed mathematical relation with a determinate extension; which respects and degrees count is set by the scientist's purpose in each case, not read off the model's structure alone. Suárez's inferential conception goes a full step further still, denying that representation requires any similarity or structural relation whatsoever: a model

represents its target just in case it is used, by a competent agent with a specific aim, to license inferences from features of the model to features of the target, a relation Suárez calls surrogative reasoning and treats as constitutive of representation rather than as evidence of some prior isomorphism.

The comparison with cartography, held to the standard the Preliminary Note set, cuts in a specific and unflattering-to-easy-analogy direction. A map's failure to be an isometry is a fact about the map and the sphere alone, provable in Chapter 4 without reference to any user's aim, purpose, or inferential practice; Tissot's theorem would hold exactly as stated if every cartographer on Earth vanished tomorrow. Suárez's inferential conception, and to a lesser degree Giere's similarity view, make representational success partly constituted by an agent's use — there is, on these views, no fact about whether a model represents its target independent of someone using it to reason with. This is the disanalogy that matters most for this volume's discipline: cartography's impossibility is purpose-independent and geometric; the semantic-view family's account of scientific representation is, on its most influential recent versions, purpose-constituted almost by definition. Where Section 3.3 asked whether the three-way distinction travels, the answer sharpens here: the distinction travels: the graded retreat from correspondence to similarity to purpose-relative use is visible in both domains, but what forces the retreat is not the same kind of fact in each, and no version of the semantic view yields, or claims to yield, a second Tissot-style theorem binding scientific models generally.

3.6 The DEKI Account: Fiction, Denotation, and What Cartography Cannot Teach It

A more recent and more articulated proposal, due to Frigg and Nguyen, analyzes scientific representation as a four-part structure they call DEKI: a model denotes its target, much as a name denotes its bearer; the model exemplifies certain properties, often via a fictional vehicle in Walton's sense — an imagined object, such as a frictionless plane or a population with no age structure, that the scientist treats as a prop in a game of make-believe governed by explicit or tacit

rules; a key translates the properties exemplified by the fictional vehicle into properties imputed to the real target, typically with an adjustment ('the target is more X than the model, in respect Y, by roughly this amount'); and imputation is the act of actually attributing the keyed-up properties to the target. The account is deliberately general enough to cover scale models, mathematical models, and diagrams within a single framework, and it is worth asking, in the spirit of Section 3.4, whether a map fits it and what the fit does and does not show.

A map fits the DEKI schema readily on its face. It denotes its territory; its symbols exemplify properties — a blue line is not literally a river but exemplifies, within the map's convention, the property 'being a watercourse'; the legend is, almost literally, Frigg and Nguyen's key, the explicit translation rule from exemplified symbol to imputed geographic property; and reading the map off against the terrain is imputation. This is not a loose fit forced onto the map from outside; the legend genuinely is a key in something close to the technical sense, and cartography may be one of the cleanest illustrations the DEKI account could ask for. But the cleanness itself reveals the limit of what the fit can show. A map's key is exhaustively codified: a legend states, completely and in advance, every symbol-to-property translation the map uses, and Chapter 4's theorem specifies, exactly and quantitatively, what the projection's own geometric key must distort and by how much at every point. The key of a scientific model built around a fictional vehicle — the frictionless plane, the perfectly rational agent, the ideal gas — is almost never codified with that completeness; whether and how far an idealization's properties transfer to the real target is typically a matter of scientific judgment exercised case by case, informed by background theory, not read off a published, invertible legend.

The honest conclusion, holding the two registers of the Preliminary Note apart, is that DEKI supplies an accurate vocabulary for describing the structure a map and a scientific idealization share — denotation, exemplification, a key, imputation — and cartography supplies an unusually complete worked instance of that structure, useful precisely because its key can be stated in closed form and its distortions proved rather than estimated. What does not transfer is any

assurance that scientific keys in general are similarly exact, similarly exhaustive, or similarly provable in their limits; nothing in Frigg and Nguyen's framework requires this, and nothing in cartography's proofs can supply it by analogy. The map is DEKI's best-behaved example, not evidence that its other examples are equally well-behaved.

3.7 The Representational Theory of Measurement: A Second Homomorphism, Independently Arrived At

Section 3.5 traced a retreat, across the semantic view of theories, from strict isomorphism toward similarity and purpose-relative use, and Section 3.6 examined a further, more articulated framework, DEKI, built to accommodate the same retreat. A literature closer to this volume's own vocabulary of correspondence, isomorphism, and homomorphism, and largely independent of the philosophy-of-science tradition Sections 3.5 and 3.6 surveyed, deserves its own mention: the representational theory of measurement, given canonical statement by Krantz, Luce, Suppes, and Tversky in their Foundations of Measurement, which asks what it is for a numerical assignment — a measurement — to represent an empirical structure faithfully. Its central technical result, the representation theorem proved for each measurement type (ordinal, extensive, conjoint, and others), establishes conditions under which a homomorphism exists from an empirical relational structure, a set of objects together with an empirically observable ordering or concatenation operation, into a numerical relational structure, the real numbers together with the ordinary ordering and arithmetic operations — homomorphism, exactly Section 3.4's term, and not accidentally: a numerical scale need not be injective in the way a strict isomorphism would require, since distinct empirical objects may be assigned the same number without falsifying the representation, precisely the pattern Section 3.4 identified in the subway map's collapse of true distance.

The convergence is worth stating plainly and then bounding just as plainly, in the spirit Section 3.5 already modeled for the semantic view. That two literatures — cartography's three-way distinction among correspondence, isomorphism,

and purpose-relative fidelity, and the representational theory of measurement's homomorphism between empirical and numerical structures — arrive independently at the same weakened relational concept is evidence that homomorphism, rather than isomorphism, is the operative notion wherever a richer structure is represented by a poorer one for a specific purpose, exactly the generalization Section 3.4 ventured cautiously and this section can now report as independently corroborated. What does not transfer is any assurance that a measurement homomorphism carries cartography's further, quantitative cost: the representation theorems of measurement theory establish existence and, via an accompanying uniqueness theorem, the permissible transformations of a given scale, but they do not, in general, specify a Tissot-style trade-off between two incompatible virtues of the representation — most representation theorems in this literature ask only whether a faithful homomorphism exists at all, not what must be sacrificed if it does. The shared word, and now the shared mathematical shape, homomorphism rather than isomorphism, travels further than Section 3.5 found the semantic view's vocabulary to travel; the quantitative trade-off of Chapter 4 still does not, and saying so here is the same discipline Chapter 11 will apply at much greater length to the rest of the philosophy of science.

3.8 Peirce's Trichotomy: Icon, Index, Symbol, and a Biographical Coincidence Not to Be Overworked

Sections 3.4 through 3.7 have tested this chapter's vocabulary of correspondence, isomorphism, homomorphism, and purpose-relative fidelity against several distinct literatures in philosophy of science; a further, older, and in one respect strikingly proximate literature is worth adding before the chapter closes: Charles Sanders Peirce's division of signs into icon, index, and symbol, developed across his writings on the logic of relatives and set out most fully in his later semiotic. A sign is, for Peirce, an icon insofar as it signifies by resembling its object in some respect independent of any actual connection between them (a portrait, a diagram); an index insofar as it signifies by an actual, often causal, connection to its object regardless of resemblance (smoke as a sign of fire, a weathervane's

position as a sign of wind direction); and a symbol insofar as it signifies by a habit or convention that could, in principle, have been otherwise (a word, a flag, the legend's stipulation that blue denotes water). Peirce is explicit, and the point matters here, that an actual sign is very rarely purely one of the three; a map is characteristically all three at once, and separating the three components sharpens rather than replaces the vocabulary Section 3.2 introduced.

A projection's iconic component is what Section 3.4's homomorphism formalizes exactly: whatever relations of adjacency, relative bearing, or relative area a given projection preserves, it preserves them by a genuine structural resemblance between source and target, the same resemblance Peirce's icon trades on informally. A map's indexical component is what Section 3.2's first notion, bare correspondence, leaves out and Peirce's category supplies: the specific causal chain — survey, instrument, satellite observation — tying a particular symbol on a particular sheet to a particular fact on the ground, the 'you are here' arrow being the purest case, an index in Peirce's full technical sense, pointing without resembling. And the map's symbolic component is exactly Section 3.6's DEKI key, made explicit and codified: the legend's stipulation that this color, this line weight, this typeface means that geographic property, a Peircean symbol in the strict sense of a sign whose connection to its object is settled by rule alone. What Section 3.2 introduced as a three-way distinction and Section 3.6 refined as denotation-exemplification-key-imputation turns out to have a third, independent formulation, arrived at for wholly different purposes a century earlier; the convergence, once again, is evidence that the tripartite structure is a genuine feature of the phenomenon and not an artifact of any one vocabulary chosen to describe it.

One further fact deserves stating carefully, precisely because of how easily it could be overstated. The author of this trichotomy is the same Charles Sanders Peirce responsible for the quincuncial projection treated in Appendix B, the conformal construction that folds the sphere, by way of an elliptic-function map, into a repeating square tile. This is a biographical fact about one unusually versatile nineteenth-century logician, not a philosophical argument, and it

licenses no inference from Peirce's semiotic to Peirce's cartography or back; the quincuncial projection's conformality is a theorem about a specific complex-analytic construction, provable without reference to any theory of signs, exactly as Section 3.5 insisted a map's failure to be an isometry holds independently of any user's purpose. What the coincidence is worth noting for is narrower and more honest: it records that the same mind found it natural to think rigorously about both how a sign of any kind manages to stand for its object and how one particular, exquisitely exact sign — a conformal map of the whole sphere — is constructed, a reminder, at the close of this chapter, that the discursive question Part I has pursued and the constructive question Part II is about to take up were never, for at least one of their historical practitioners, two different subjects.

3.9 Diagrammatic Reasoning: Peirce, Stjernfelt, and the Map as an Argument

Section 3.8 introduced Peirce's trichotomy of icon, index, and symbol and placed the map, with qualifications, among the icons. Peirce made a further claim specifically about a subclass of icons he called diagrams: unlike a portrait, which merely resembles its object, a diagram exhibits relations among its parts that mirror relations among the parts of what it represents closely enough that operating on the diagram — extending a line, comparing two lengths, tracing a path — licenses a genuine inference about the represented relations, not merely a further description of the picture.

Frederik Stjernfelt's extended treatment of this idea, diagrammatology, generalizes the point beyond Peirce's own geometric examples: any sufficiently structure-preserving representation supports what Stjernfelt calls diagrammatical reasoning, inference performed by manipulating the representation itself rather than by manipulating propositions about it. Reading a map for the shortest route, or comparing two settlements' apparent distance from a coastline, is diagrammatic reasoning in exactly this sense — an operation

on the sheet that is trusted to answer a question about the territory without any separate proof being consulted.

The trust is only as good as what Part II shows a projection actually preserves, and this is where Section 3.9 hands off cleanly to Chapter 4. Diagrammatic reasoning performed on a non-conformal projection by comparing apparent angles, or on a non-equal-area projection by comparing apparent regions, is not a weaker form of inference than reasoning on a conformal or equal-area one; it is, in the relevant respect, invalid reasoning, since the diagram no longer mirrors the relation being read off it. Peirce's insight that a map licenses inference, not just depiction, turns out to sharpen rather than soften the stakes of Tissot's theorem: the theorem specifies, chapter and verse, exactly which diagrammatic inferences a given projection can support and which it silently invalidates.

3.10 Wittgenstein's Picture Theory and the Limits of the Map-as-Proposition Analogy

Section 3.2's discussion of isomorphism and structure-preservation brushed against, without naming, one of the twentieth century's most influential representational theories: the early Wittgenstein's picture theory of the proposition, developed in the Tractatus Logico-Philosophicus, on which a meaningful proposition is a picture (Bild) of a possible state of affairs, sharing with it a common logical form even though the proposition's elements (words) need not resemble the state of affairs' elements (objects) in any sensuous way.

The theory's most distinctive and most relevant claim for this volume is that a picture cannot depict its own pictorial form: a picture can show that things are thus-and-so, but it cannot show what makes it a picture of anything at all, the form it shares with reality, since to depict that form the picture would have to step outside itself. A map exhibits the same limit concretely — no amount of detail on a sheet can show, on the sheet, what makes marks on paper a representation of a coastline rather than an arbitrary pattern; that fact belongs to the map's relation to its territory, not to anything printable within its borders, a

point Section 3.6's treatment of the map's fictional, DEKI-theoretic status already required without invoking Wittgenstein's name for it.

The analogy has a clear edge, consistent with the Preliminary Note's discipline. Wittgenstein's own later work abandoned the picture theory's generality, replacing it with the plural, rule-governed language games of the Philosophical Investigations; and the picture theory was in any case a thesis about propositions in general, a binary claim about what can and cannot be shown, not a quantitative one. Chapter 4's theorem is exactly the quantitative claim the picture theory does not supply: not merely that a map cannot show its own form of depiction, but a precise, provable account of what it must additionally sacrifice, and by how much, in depicting anything else.

3.11 A Coda: What the Test Case Has Established, and What It Has Deliberately Left Open

Sections 3.1 through 3.10 have used the map to test, rather than merely illustrate, several competing accounts of what a representation is: strict isomorphism, the weaker homomorphism the representational theory of measurement actually requires, Peirce's icon–index–symbol trichotomy, and the DEKI account's attempt to say what a map denotes when, strictly, no perfect correspondence exists to denote it by. What the test case has established, cumulatively, is narrow but load-bearing: that 'representation' is not one relation but a family of related, formally distinguishable relations, that cartography instantiates several members of the family at once and at different levels of its own construction, and that a philosophical account adequate to the map must be precise about which member of the family it is claiming to describe, on pain of the conflations Section 3.2 catalogued at the outset.

What this coda does not claim, and flags here rather than leaving to be assumed by momentum into Chapter 4, is that any of this settles the question of representation for domains structurally unlike cartography. The map is a good test case exactly because it has an independently characterizable target, the

curved surface Gauss's theorem constrains, against which formal notions of correspondence can be checked with unusual precision; most representations, in science and elsewhere, do not have targets this well behaved, and Chapter 11 will return, much later in the volume, to ask exactly how much of what has been established here survives that loss of good behavior. For now, Part II turns from the general question of what a map represents to the specific, provable question of what it costs.

PART II

Impossibility as a Philosophical Figure

Chapter 4 — Tissot's Theorem and the Price of Representation

4.1 Stating the result without re-proving it

The technical content of this chapter is established, with full proof, in the companion volume The Impossibility Theorems of Cartography (Chapters 2–4), which derives Tissot's indicatrix from the local linearization of a projection and proves, from it, that no projection of a curved surface onto the plane can be simultaneously conformal (angle-preserving) and equal-area (area-preserving) except at isolated points. This chapter does not repeat that derivation. It takes the result as given and asks what it means to have proved, rather than merely observed, that a specific pair of representational virtues cannot both be had in full.

The result has a precise shape worth stating carefully before any philosophical use is made of it. At any point of any projection, the local distortion is characterized by two principal scale factors, conventionally written a and b, describing how much lengths are stretched in two perpendicular directions. The projection is conformal at that point exactly when a equals b (the local picture is a uniformly scaled, undistorted copy of a small neighborhood); it is equal-area at that point exactly when the product ab equals 1 (areas are preserved even though shapes may be sheared). A point where both conditions hold simultaneously requires a = b = 1: no distortion at all, of either kind. Tissot's theorem, and the deeper Gauss–Bonnet argument behind it, show that this can hold at isolated points or along isolated lines, but never — for a genuinely curved region — everywhere at once. The trade-off is not a rule of thumb; it is a theorem about what values a and b can jointly take, given the intrinsic curvature of the surface being mapped.

4.2 A trade-off, not a hierarchy

The philosophically important feature of Tissot's result is not merely that some distortion is unavoidable — a claim already drawn from Chapter 1 — but that the distortion is distributed across two independent axes (angular and areal fidelity) that trade off against one another rather than being ranked on a single scale of 'more' or 'less' accurate. A conformal projection is not a more accurate projection than an equal-area one; it is accurate along a different axis, at the cost of the axis the other preserves. This matters philosophically because it blocks a natural but mistaken picture of representational failure as a single quantity — a general-purpose 'distortion' that a sufficiently good map simply has less of. There is no such single quantity for projections in general. There are (at least) two, and improving one, past a certain point, provably worsens the other.

This structure — not one gradient of fidelity but several, non-collapsible into one, with an exact trade-off law linking them — is the specific, non-metaphorical content this volume proposes to carry forward. Whether an analogous plurality of independent, non-collapsible virtues holds for other representational practices is, again, a question this volume can raise but not settle by fiat; the cartographic case is offered as a rigorously worked instance of the general shape such a plurality can take, not as proof that it always does.

4.3 Choosing among the survivors

If no projection can have both virtues, the question a cartographer actually faces is not 'how do I avoid distortion' but 'which distortion do I choose, for which purpose, and at whose expense.' A navigational chart (Mercator) chooses conformality, because a sailor needs the bearing read off the chart to be the true compass bearing, and is willing to accept wildly exaggerated areas at high latitudes as the price. A thematic map of population density chooses equal-area representation, because comparing the size of shaded regions only means something if area is preserved, and is willing to accept sheared, angularly distorted shapes as the price. Neither choice is a failure to achieve the other virtue through insufficient skill; each is the correct response to a specification that determines, once stated, which theorem-governed trade-off point is appropriate.

The next chapter turns from this pointwise, quantitative trade-off to a different and stricter kind of impossibility — one that admits no trade-off at all, because the obstruction is topological rather than metric.

4.4 The Named Projections as Case Studies

Sections 4.2 and 4.3 stated the trade-off Tissot's theorem forces in general terms: no projection is both conformal and equal-area, and the choice among the survivors is a choice of purpose, not a search for a nonexistent ideal. The abstraction is easier to hold onto with three concrete cases in view, each a real, historically significant choice made within exactly the space the theorem describes. The Mercator projection is conformal: angles are preserved everywhere, which is precisely why a straight line on a Mercator chart is a line of constant compass bearing, the property that made it indispensable to sailing navigation from the sixteenth century onward. The cost, exactly as Tissot's theorem requires, is area: Mercator's projection inflates regions with distance from the equator without bound, a consequence returned to, in its political dimension, in Chapter 13. The Albers equal-area conic projection sits at the theorem's other pole: it preserves area everywhere at the cost of angle, making it the standard choice for thematic maps, population density, land use, resource distribution, where a reader comparing the apparent size of two regions needs that comparison to be true, and where the shape distortion this accuracy costs is a price most such uses can afford. Between these two lies a family of compromise projections, of which Robinson's, adopted for a period by the National Geographic Society for general-purpose world maps, is the best known: neither conformal nor equal-area, it distorts both angle and area everywhere, deliberately, in exchange for distortions that are individually smaller and, its designer argued, visually less alarming than either extreme.

What these three cases make vivid is a feature of the trade-off that the abstract statement in 4.2 can leave too easy to accept without friction: the theorem does not merely permit a compromise projection like Robinson's, it guarantees that the compromise, too, pays a cost, at every point, of a form the theorem itself specifies

exactly. There is no projection, however cleverly designed, that finds a loophole in Tissot's indicatrix; Robinson's projection is not a fourth kind of thing exempt from the theorem but a particular, named point in the same constrained space Mercator's and Albers's projections also occupy, chosen, as Section 4.3 already argued, for its purpose, and paying, as every point in that space must, a determinate price for it.

4.5 A Fourth Case: The Dymaxion Projection and the Limits of Rearranging the Cost

Sections 4.2 through 4.4 traced the trade-off through three named cases — Mercator, Albers, Robinson — chosen because each occupies a different, clearly labeled point in the space Tissot's theorem describes. A fourth case is worth adding, not because it changes the shape of the trade-off but because it tests, in an unusually pointed way, the limits of what rearranging the cost can buy. Buckminster Fuller's Dymaxion projection, unfolded onto the faces of an icosahedron before being laid flat, distributes its distortion across the seams between the icosahedron's twenty triangular faces rather than concentrating it, as Mercator's does, at the poles; the result is a projection under which no single landmass suffers the extreme, visually dramatic distortion Mercator's high latitudes do, and continental shapes and relative sizes read, to an untrained eye, as considerably more faithful than either Mercator's or even Robinson's compromise. This is a genuine achievement within the space Section 4.2 described, and it is worth being precise about what kind of achievement it is. The Dymaxion projection does not evade Tissot's theorem — no projection can — it instead exploits the theorem's silence about where distortion is placed, choosing an unfolding that scatters the unavoidable cost across many small, hidden seams rather than paying it in one large, visible deformation.

This is the trade-off's residual freedom, in Chapter 6's sense, exercised with unusual ingenuity, and it is worth marking precisely why ingenuity here does not amount to a fifth kind of thing exempt from the classification Section 4.4 already gave. A cartographer choosing where to place a projection's inevitable

seams is choosing a distribution of cost, not its magnitude; the total distortion integrated over the sphere is not reduced by clever unfolding, only redistributed into a pattern the human eye finds less alarming, for reasons having more to do with the visual psychology of large contiguous deformation than with any quantity Tissot's theorem itself governs. The Dymaxion projection is, in this precise sense, evidence for Section 4.4's claim rather than an exception to it: even the most inventive available point in the trade-off space still pays, at every point, the exact price the theorem specifies, and what ingenuity achieves is a better answer to the different, non-geometric question Chapter 6 will pose next — not how much must be paid, but where, and to whose benefit, the payment is best made.

4.6 Tissot’s Trade-off as a Pareto Frontier

Sections 4.2 and 4.3 described the space of possible projections as admitting a trade-off between angle-preservation and area-preservation, with no projection achieving both and every projection occupying some point along the space of achievable compromises; a vocabulary from outside cartography names this structure exactly rather than merely by analogy, and is worth introducing before the chapter closes. Vilfredo Pareto’s concept of efficiency, developed for problems of resource allocation and now standard across economics, engineering, and multi-objective optimization generally, calls an outcome Pareto-efficient if no alternative exists that improves at least one of the outcome’s specified objectives without making some other objective worse; the set of all Pareto-efficient outcomes, plotted against the objectives being traded off, traces a Pareto frontier, a boundary beyond which the objectives cannot jointly be pushed, though movement along the boundary itself, trading one objective’s value for another’s, remains freely available.

This is not an analogy loosely applied to Chapter 4’s result; it is the same mathematical structure under a different name, and saying so precisely is exactly the discipline the Preliminary Note asked for. Take angle-fidelity and area-fidelity, each measured at a point by some agreed deviation from the ideal

Tissot indicatrix, as the two objectives; a projection that could be modified to improve one without worsening the other anywhere would be, in the ordinary sense, an improvement no cartographer would have reason to decline, and Tissot's theorem is exactly the claim that no projection sits strictly inside the frontier this way — every projection already lies on the boundary the trade-off allows, for the simple reason that conformal projections (Mercator among them) already achieve one objective's maximum absolutely, and equal-area projections (Albers, and the Lambert azimuthal equal-area case) already achieve the other's, with every intermediate choice, Robinson's among them, occupying a point along the frontier connecting the two extremes rather than a point strictly dominated by some undiscovered alternative.

What this reframing adds to Section 4.3's already-stated conclusion, that purpose disciplines the choice among survivors, is a sharper diagnosis of what kind of choice that is. Multi-objective optimization's own literature distinguishes, as a matter of course, between finding the Pareto frontier, a fact about the problem's structure, solvable in principle without reference to anyone's preferences, and selecting a point on it, which requires a scalarization, an explicit or implicit weighting of the objectives against each other supplied from outside the optimization problem itself. Chapters 4 and 6 have already argued for exactly this two-stage structure without the vocabulary: Tissot's theorem finds the frontier, and Poincaré's conventionalism, examined in Chapter 6, is this volume's account of where the scalarizing weights, angle over area for a sailor, area over angle for a demographer, actually come from. The frontier itself asks nothing of anyone's purpose; the point chosen on it asks for nothing else, and naming the two halves in optimization theory's own vocabulary makes visible, in a single well-worked technical discipline outside cartography, exactly the division of labor Chapters 4 through 6 have been arguing for by other means.

4.7 Compromise Projections and the Ordinary Cartographer's Practical Criterion

Sections 4.4 through 4.6 treated named projections as case studies positioned on or near the Pareto frontier of Tissot's trade-off — each optimal, in its region, along some single virtue or explicit combination of virtues. A further and in practice more common category was deliberately set aside until now: compromise projections, such as Robinson's or the Winkel Tripel, which are neither conformal, nor equal-area, nor equidistant, and are not derived by optimizing any explicit combination of those quantities at all.

Arthur Robinson's own account of constructing his projection is instructive precisely because it is not a proof: he tabulated coordinates by iterative visual adjustment, judging by eye which of several trial renderings looked least distorted across the whole globe, rather than by minimizing any of the error metrics Section 4.6's Pareto frontier is built from. The result sits inside the frontier, not on it — dominated, on any single named criterion, by a specialist projection built for that criterion — and was nonetheless adopted by the National Geographic Society for general reference maps in preference to more rigorously optimal alternatives, including the Mercator projection it replaced for that purpose.

The philosophical point is not that Robinson's method was sloppy where a better one was available, but that a fourth, non-geometric criterion was operating openly beside the three Tissot-theoretic ones: legibility and visual plausibility to an untrained reader, which no combination of conformality, area-preservation, and equidistance guarantees or even correlates with in general. This anticipates Chapter 6's conventionalism a step early: even after the trade-off's mathematical shape is completely fixed by the theorem, and before any dispute about which named virtue to prioritize arises, a further judgment — aesthetic, pedagogical, frankly a matter of what a reader's eye will accept — is required to pick the actual working point, and this judgment is no more forced by the geometry than the choices Section 6.3 will call constrained without being forced.

4.8 Chebyshev's Theorem: An Exact Optimum for Small Conformal Maps

Sections 4.4 through 4.7 treated named projections and compromise projections as points inside or on Section 4.6's Pareto frontier, evaluated globally, across the whole sphere. A narrower but exactly solved question sits inside that larger picture: for a small, bounded region rather than the whole globe, is there a single best conformal map, and can 'best' be made precise rather than judged by eye as Section 4.7's Robinson projection was?

Pafnuty Chebyshev answered this in 1856 for the restricted but well-posed case of minimizing the maximum scale distortion of a conformal projection of a given bounded region: the optimal conformal map, among all conformal maps of that region, is the one whose scale factor is constant along the region's boundary, a result later given a complete proof by Grave in 1896. Unlike Chapter 4's central theorem, which forbids a combination of virtues outright, Chebyshev's result operates entirely within the conformal branch of that trade-off, already having conceded area and shape distortion elsewhere, and asks the narrower question of how to minimize scale variation specifically, given that concession.

The result is a useful corrective to any impression Sections 4.4–4.7 might have left that everything below the level of Tissot's theorem is judgment and compromise. Some sub-questions inside the trade-off are themselves fully solved: given that a projection will be conformal, and given a fixed region, there is a demonstrably optimal choice, not a matter of taste. The lesson generalizes usefully to Section 6.3's claim that convention operates only where genuine freedom remains — Chebyshev's theorem marks one more place, this time inside the conformal branch rather than at the top level, where that freedom turns out, on inspection, not to be as wide as it first appears.

4.9 Rate–Distortion Theory: A Structural Cousin from Information Theory

Section 4.6 named the trade-off Tissot's theorem forces as a Pareto frontier, borrowing the vocabulary of multi-objective optimization to say precisely what kind of structure links angular and areal fidelity. A second borrowed vocabulary,

from a field with no more interest in cartography than economics has, names the same shape from a different direction and is worth setting beside it before the chapter closes. Claude Shannon's rate-distortion theory, developed in 1959 to answer a question about compressing a signal rather than projecting a sphere, proves that for any source and any tolerated level of distortion D, there is a minimum achievable rate R(D), bits per symbol, below which no encoding scheme, however cleverly designed, can represent the source and stay within that distortion — and that R(D) is a strictly decreasing function of D wherever D is not already so generous that no compression is needed at all. Lower distortion is achievable, but only by spending more rate; the two quantities trade off against a fixed curve, not against each other by rule of thumb, in exactly the sense Section 4.6's frontier fixes the trade between angular and areal fidelity.

The parallel is worth stating with the same discipline Section 4.6 insisted on for the Pareto vocabulary: it is a structural analogy, not a derivation of one theorem from the other, and the two results concern different mathematical objects — a source distribution and an encoding, on one side; a Riemannian metric and a map to the plane, on the other. What they share is the general shape this volume has already isolated as its recurring find: two quantities, each independently desirable, jointly bounded by a proven curve rather than by an ad hoc rule, so that a claim to have improved one without cost to the other is not a report of ingenuity but a report of an error, checkable against the curve itself. Where Section 4.6 asked the reader to see Tissot's theorem as one instance of a Pareto frontier, this section asks the same reader to notice that engineering, faced with an entirely unrelated problem — how much a signal can be compressed before something is unrecoverably lost — proved a theorem of the identical shape, decades before either field had occasion to compare notes with the other.

One disanalogy is worth flagging rather than smoothing over, in the same spirit Section 5.6 later insists on for the hairy ball theorem's scope. Rate-distortion theory's curve depends on the statistics of the particular source being compressed — a different signal yields a different R(D) — where Tissot's trade-off depends only on the sphere's curvature, fixed once and for all,

independent of which projection is attempted. The cartographic frontier is, in this sense, the more rigid of the two: geodesy supplies one curve for all mapmakers, where information theory supplies a family of curves, one per source. What survives the disanalogy, and is the only claim this section makes, is the shared logical form — a proven, non-negotiable boundary between two desiderata — not a claim that cartography and data compression are otherwise the same kind of problem.

4.10 The Winkel Tripel and the National Geographic Society: A Convention Adopted, Not Derived

Section 4.7 described the compromise projection as a practical criterion, distinct from the mathematically privileged conformal and equal-area extremes, by which working cartographers navigate the trade-off Tissot's theorem proves is unavoidable. A single, well-documented institutional episode makes the point concretely rather than abstractly. In 1998 the National Geographic Society, after decades of using the Robinson projection, another compromise, adopted the Winkel tripel projection, devised by Oswald Winkel in 1921 as an arithmetic mean of the equirectangular and Aitoff projections, as its standard for world maps. Neither projection is forced by any theorem this volume has proved or will prove; both sit inside the survivor set Section 4.3 described, minimizing some blend of angular and areal error rather than eliminating either, and the choice between them, and the further choice to switch from one to the other, was made by an editorial and cartographic staff weighing legibility, aesthetic convention, and continuity with prior publications, not by a calculation with a unique correct output.

The episode is worth setting beside Section 4.9's discussion of rate–distortion theory for a reason specific to this volume's discipline about registers. Rate–distortion theory shows that trading two kinds of cost against each other along an optimal frontier is a structural feature reappearing wherever a lossy encoding is required, cartographic or otherwise; the Winkel tripel adoption shows what it looks like when a real institution, faced with the same frontier

Chapter 4 has proved exists, must additionally decide where on that frontier to stand, without the geometry itself supplying the answer. The National Geographic Society's statement accompanying the 1998 change cited the Winkel tripel's 'flattened' polar regions and reduced distortion at high latitudes relative to the Robinson projection as the deciding consideration for its particular use case, world reference maps intended for a general audience — a criterion external to Tissot's theorem, resolvable only by knowing what the map was for, exactly as Section 4.3 anticipated in the abstract.

This is also, not incidentally, a first concrete instance of the theme Chapter 6 will take up as its central subject: that the choice among mathematically permissible projections is constrained by the theorem without being determined by it, leaving a residue of genuine, non-arbitrary choice that Poincaré's conventionalism, rather than the geometry of the sphere, is the right framework for understanding. The Winkel tripel did not become the National Geographic Society's standard because a proof singled it out from among Chapter 4's survivors; it became the standard because an institution, accountable to readers and to its own history of prior maps, decided it was the better tool for a purpose no theorem specifies. Chapter 6 returns to this structure in full generality; here it is enough to have shown that the abstract trade-off of Sections 4.1 through 4.9 has a concrete, datable, institutional history exactly where the theorem predicts one must exist.

4.11 A Second Institutional Episode: The Equal Earth Projection and the Controversy It Was Built to Answer

Section 4.10 traced one institutional resolution of the trade-off Tissot's theorem proves unavoidable, the National Geographic Society's 1998 adoption of the Winkel tripel; a second, more recent episode belongs beside it, arising from a controversy Section 13.7 examines in its own right when the volume turns from the trade-off's existence to its political uses. In March 2017 the Boston Public Schools adopted the Gall–Peters projection, an equal-area projection Section 13.7 discusses at length, as the standard for its classroom world maps, explicitly to

correct what the district described as a long-standing distortion of relative country sizes associated with Mercator-derived world maps. The switch drew renewed attention to a familiar complaint about the Gall–Peters projection itself: that its equal-area property is purchased, exactly as Chapter 4's trade-off predicts, at the cost of severe shape distortion, elongating landmasses vertically near the poles and compressing them near the equator in a manner many cartographers had long argued undermined the map's legibility and visual appeal even as it corrected the area problem.

In direct response, three cartographers, Tom Patterson, then of the United States National Park Service, Bernhard Jenny of Monash University, and Bojan Šavrič of Esri, set out to find an equal-area projection that preserved correct relative country sizes without the Gall–Peters projection's visual costs, and, finding no existing option that met their aesthetic criteria, developed a new one, published in August 2018 in the International Journal of Geographical Information Science under the name Equal Earth. The projection sits, in the vocabulary Section 4.3 introduced, squarely among the equal-area survivors of Tissot's theorem, sacrificing conformality exactly as the Gall–Peters projection does, while differing from it in the further, non-forced choice of a pseudocylindrical form inspired by the Robinson projection's proportions; it was adopted within months by NASA's Goddard Institute for Space Studies for its own global temperature maps and has since become a standard option in widely used geographic information software. Like the Winkel tripel's adoption by the National Geographic Society, the Equal Earth projection's creation and rapid uptake was not required by any theorem; what Chapter 4's proof forced was only that some cost be paid for the areal accuracy the Boston Public Schools wanted, and Section 4.3's survivor set left the particular shape of that cost, and the particular map eventually chosen to pay it, for institutions and individual cartographers, not geometry, to decide.

Chapter 5 — The Hairy Ball and the Limits of Global Orientation

5.1 A different register of impossibility

Every impossibility examined in the previous chapter is, in a precise sense, negotiable: conformality and equal-area representation cannot both be had everywhere, but a projection can trade a controlled, quantified amount of one for a controlled, quantified amount of the other, and a cartographer can choose the trade-off point that suits a given purpose. This chapter treats a different kind of result, established in the companion volume (The Impossibility Theorems of Cartography, Chapter 6) from the hairy ball theorem of algebraic topology: a continuous, nowhere-vanishing tangent vector field on the two-dimensional sphere does not exist. Nothing about this obstruction is negotiable. There is no partial version of it, no trade-off curve, no choice of parameter that buys back some of what is lost. Either the required field exists or it does not, and on the sphere it does not, for every attempted construction, without exception.

The proof, given in full in the companion volume via the degree of the Gauss map (and independently via a Poincaré–Hopf index argument, developed further in the same volume's later chapters on Goldberg polyhedra), is not repeated here. What matters for this chapter is the shape of the result: it is a statement about the entire sphere at once, not about any single point or any local neighborhood. A vector field can be built that is smooth and non-vanishing on all of the sphere except one point, or two, or any finite set — the obstruction is not a difficulty distributed evenly across the surface, the way angular and areal distortion are; it is concentrated, forced by the topology, into at least one singularity that cannot be smoothed away, wherever the construction is attempted.

5.2 The cartographic instance: no universal field of fixed directions

Chapter 6 of the companion volume applies the theorem directly to cartography: a globally continuous field of directions pointing toward a fixed reference point C — the bearing, at every point P of the sphere, of the great-circle direction back to C — cannot be defined without a singularity. At C itself the direction is

undefined by construction (a point has no bearing to itself), and at the antipode of C every direction is, in the relevant limiting sense, equally a direction toward C, so the field cannot be extended continuously across that second point either. Two singularities, forced, not by any deficiency of technique but by the theorem: this is the topological content behind the retroazimuthal problem examined at length, historically and mathematically, in the companion volume The Retroazimuthal Problem, and treated in this volume's own terms in Part III below.

The distinction from Chapter 4 is worth making sharp. Tissot's obstruction says: you may have angular fidelity everywhere, or areal fidelity everywhere, but not both — and if you want a bounded compromise, you may have that too, continuously tunable, at a quantified cost. The hairy ball obstruction says something categorically different: you may not have a global, singularity-free field of directions toward a fixed point at all, on any terms, at any cost — the singularities are not a cost to be paid but a topological necessity to be accommodated. The three historical solutions to the retroazimuthal problem (Littrow, Craig, Hammer) do not defeat this obstruction; each accommodates it in a different way, either by giving up conformality, by giving up global injectivity, or by giving up freedom in the choice of the reference point C — precisely the menu of concessions a genuine impossibility, rather than a mere difficulty, forces.

5.3 Metaphor and argument, kept apart

The phrase 'you can't comb a hairy ball flat' has itself become a piece of loose popular mathematics, often reached for as an all-purpose image for the impossibility of achieving global order, coherence, or consensus in some human domain — political, social, psychological. The Preliminary Note to this volume promised to resist exactly this move without argument, and the promise is worth honoring explicitly here, at the point where the temptation is strongest. The theorem is a statement about continuous tangent vector fields on a specific topological space, the two-sphere; it says nothing, by itself, about committees,

opinions, or selves. Where this volume later asks (Part III, and again in the closing chapters of Part IV) whether the structure the theorem exhibits — a global requirement that provably cannot be met without at least one forced exception — recurs in phenomenological or social domains, it will do so only by identifying an actual topological or structural analogue of the sphere and the vector field in question, not by re-using the theorem's vivid image on the strength of a family resemblance. Short of that work, the theorem licenses nothing beyond cartography, however tempting the sentence 'no map of directions is ever complete' sounds when detached from the surface it is a theorem about.

5.4 What remains for Part III

Chapter 4 established that partial fidelity is negotiable; this chapter has established that a specific kind of global fidelity — a complete, singularity-free field of bearings toward a fixed point — is not negotiable at all, only accommodated. Part III turns from this abstract, third-person result to the first-person question it bears on most directly: what it is, concretely, to stand at some point P and seek the direction of a fixed center C one does not occupy — a structure common to prayer direction, to homesickness, and to every act of orienting oneself toward an elsewhere. The chapter that opens Part III (Chapter 7) takes up that question through the retroazimuthal problem's most sustained historical instance, the determination of the qibla.

5.5 Why Two, Not One or None: Parity as a Hidden Theorem

Section 5.4 followed the companion volume in noting that the retroazimuthal field is forced to fail at two points, the center C and its antipode, and Chapter 7 will make a great deal, philosophically, of the difference between these two singularities. It is worth pausing on a question the number two itself raises and the hairy ball theorem, stated only as a nonexistence result, does not by itself answer: why two failures rather than one, or twenty? The answer is a further,

and so far unstated, piece of the same topology, a consequence of the Poincare-Hopf theorem, of which the hairy ball theorem is the sphere's special case. Poincare-Hopf does not merely forbid a nonvanishing vector field on the sphere; it specifies exactly how much vanishing is unavoidable, by requiring that the sum of the indices of a vector field's zeros equal the surface's Euler characteristic, two, for the sphere. A single zero of index two satisfies this exactly, and is realizable, but is topologically fragile in a sense a philosophical reader can grasp without the surrounding machinery: it is a degenerate configuration, unstable under the smallest perturbation of the field, which generically splits a single index-two zero into two index-one zeros. The retroazimuthal field's actual behavior, a well-behaved singularity at C and a second, differently-behaved singularity at the antipode, discussed at length in Chapter 7, is exactly this generic, stable configuration, not a special or unusual case.

The philosophical point available here anticipates Chapter 7's distinction between the two singularities' character without collapsing into it: the theorem does not only guarantee that the orienting project must fail somewhere, it constrains, quantitatively, how the failure is distributed, and the case this volume has followed throughout, one failure at arrival, a structurally different failure at the point of maximal distance, is the topologically typical case, not a special feature of the qibla or of any other single historical instance of the retroazimuthal problem. Wherever a fixed center is to be faced from every other point on a sphere, this same parity is waiting to be found, whether or not the particular culture asking the question ever had occasion to notice its second, antipodal half.

5.6 Beyond the Two-Sphere: Odd Dimensions and the Specificity of the Obstruction

Chapter 5's theorem was stated, deliberately, only for the two-sphere, because that is the surface cartography needs; it is worth asking, briefly, whether the obstruction is a general fact about spheres or a specific fact about two dimensions, because the answer sharpens what the theorem does and does not

license even within pure topology, before any cartographic or phenomenological use is made of it at all. The Poincare-Hopf theorem invoked in Section 5.5 answers the question completely: a smooth, nowhere-vanishing tangent vector field exists on the n-sphere if and only if its Euler characteristic vanishes, and the Euler characteristic of the n-sphere is 1 + (-1)^n, which is zero exactly when n is odd and 2 when n is even. The two-sphere, n = 2, is even, and so is every ordinary globe a cartographer might map; but the circle (n = 1) and the ordinary three-dimensional sphere embedded in four-dimensional space (n = 3) are odd, and both admit smooth, nowhere-vanishing tangent vector fields — a fact easy to verify directly for the circle, where the field 'always point clockwise' never vanishes, and provable for the three-sphere by treating its points as unit quaternions and using quaternion multiplication to rotate a fixed tangent direction consistently at every point, a construction with no analogue on the ordinary globe because the two-sphere supports no such multiplicative structure making every point behave alike.

The philosophical use of this fact, held to the same discipline Section 5.3 insisted on, is narrow but genuine: it forecloses, in advance, a tempting but false generalization of the retroazimuthal obstruction — the thought that 'no sphere can support a global field of fixed directions' as though sphere-ness alone were doing the forcing. It is not sphere-ness; it is the specific, even-dimensional sphere's topology, and the fact that the obstruction disappears for the circle and for the three-sphere is direct evidence that the theorem is tracking a precise structural fact about the surface in question rather than a vague intuition about curved things in general — precisely the standard of precision the Preliminary Note asked this volume to hold itself to throughout, and a small but clean illustration, internal to the mathematics itself before any philosophical extension is attempted, of how much is lost when a theorem's exact scope is replaced by an evocative paraphrase of its conclusion.

5.7 A Physical Instance, Not a Metaphor: Nematic Defects and the Same Theorem

Section 5.3 insisted on keeping metaphor and argument apart, and the caution is worth restating before adding one further case, because what follows is not a metaphorical extension of the hairy ball theorem but a second, literal instantiation of it, in a domain with no cartographic content at all. A nematic liquid crystal is a material whose constituent molecules, rod-shaped, align locally along a common axis, described mathematically by a headless unit vector, or more precisely by a line field rather than a true vector field, at every point of the material; when such a material is confined to a spherical droplet, a configuration studied extensively in soft-matter physics, the alignment field is defined on the surface of a sphere in essentially the same sense Chapter 5's tangent vector field was, and is subject to the identical topological obstruction Chapter 5 proves: no smooth, defect-free alignment of the molecules can cover the whole sphere, because the sphere's Euler characteristic, the same invariant Section 5.5 used to explain why exactly two singularities are forced rather than one or none, forbids it.

The physical consequence is directly observable and has a name in the condensed-matter literature independent of any cartographic vocabulary: topological defects, points where the local alignment becomes undefined or singular, are not a flaw in any particular experimental preparation but a forced feature of nematic order on a sphere, and their minimum total topological charge is fixed by the same Poincaré–Hopf accounting Section 5.5 already introduced. A spherical nematic shell characteristically settles, at equilibrium, into a configuration with four half-strength defects arranged tetrahedrally, or into other charge-conserving arrangements, but never into a defect-free configuration, for exactly the reason no continuous field of directions toward a fixed point can cover the whole sphere without at least the two singularities Chapter 5's cartographic case requires. The specific number and arrangement differ, because a nematic's headless line field is not the same mathematical object as the retroazimuthal problem's vector field and satisfies a different, related index-counting rule, but the underlying obstruction, that some singularity is unavoidable on a sphere, is the identical theorem, not a cousin of it.

This case is worth including for a reason distinct from Section 5.6's extension to odd dimensions: that section stayed inside mathematics, generalizing the theorem's statement; this one leaves mathematics entirely and finds the same statement forced on a physical system with no interest in cartography, orientation, or representation of any kind, by molecules that arrange themselves according to free-energy minimization rather than by anyone's theorem. The philosophical use of this fact is narrow and exactly bounded, in the spirit Section 5.3 asked for throughout: it is evidence that the hairy ball obstruction is a fact about spheres, full stop, indifferent to whether the field living on the sphere is a cartographer's bearing, a physicist's molecular alignment, or a meteorologist's surface wind, which cannot, for the same reason, ever be entirely calm everywhere at once. It licenses no further claim about orientation, dwelling, or the qibla beyond what Chapter 7 will independently establish; it only confirms, from a domain that owes cartography nothing, that Chapter 5's theorem describes the sphere rather than any human practice performed upon it.

5.8 Arrow's Impossibility Theorem: A Structural Cousin from Social Choice

Chapter 5 proved a topological impossibility: no continuous, everywhere-defined field of directions toward a fixed point can exist on the sphere. It is worth asking whether the argumentative form of that result — a short list of individually reasonable conditions shown, by direct construction of a contradiction, to be jointly unsatisfiable — recurs outside geometry, and a genuine, independently discovered instance is Kenneth Arrow's theorem in social choice: no procedure for aggregating three or more individuals' rankings of three or more options into a single collective ranking can simultaneously satisfy unrestricted domain, the Pareto principle, independence of irrelevant alternatives, and non-dictatorship.

The two proofs share nothing at the level of mathematical content — one is a statement about continuous vector fields on a compact surface, the other about finite combinatorics of orderings — but they share a shape worth naming explicitly, since Section 5.3 has already insisted on keeping shared shape apart

from shared substance. Both results close off, permanently and by proof rather than by exhaustive search, the response that a cleverer construction might yet satisfy every condition at once; both were, in their own fields, initially met with attempts to find the loophole the proof had already foreclosed, exactly as Section 5.1 recounts happened, more than once historically, with candidate global wind-direction fields.

What the parallel is good for, kept at the level of form the Preliminary Note requires: it shows that Part II's recurring finding — that some configurations of jointly desirable properties are not merely difficult to achieve but strictly forbidden — is not a peculiarity of curved-surface geometry. It is a shape that recurs wherever a single global output (a vector field, a social ranking) is demanded from purely local or pairwise conditions that look individually innocuous. Cartography and social choice, developed in total independence of one another, converge on the same warning to any designer tempted to add one more apparently reasonable requirement without first checking whether the whole list remains jointly satisfiable.

5.9 Borsuk–Ulam: A Second Topological Cousin, Closer to Home Than Arrow

Section 5.8 introduced Arrow's theorem as a structural cousin to the hairy ball theorem, sharing an argumentative form but no mathematical content. A second cousin exists that shares rather more: the Borsuk–Ulam theorem, which states that any continuous function from the surface of a sphere to the plane must assign the same value to at least one pair of antipodal points — stated, unlike Arrow's theorem, in the same topological universe as Chapter 5's own result, about continuous functions on spheres.

The standard illustration is meteorological rather than cartographic, and worth stating for its own vividness: at any moment, there exists at least one pair of points on opposite sides of the Earth with exactly the same temperature and exactly the same barometric pressure simultaneously, not as an empirical

regularity but as a mathematical certainty following from Borsuk–Ulam applied to the pair of continuous scalar fields temperature and pressure. This is a genuinely different obstruction from Chapter 5's: the hairy ball theorem concerns a vector field (a direction at every point) and forces a zero, a point of total failure; Borsuk–Ulam concerns a pair of scalar fields (numbers at every point) and forces a coincidence, two points in perfect if accidental-seeming agreement, not a breakdown at either.

Both theorems nonetheless belong to the same technical family as instances of what topologists classify together as fixed-point and coincidence results on spheres, proved by closely related methods, which is why this cousin is closer to Chapter 5's own result than Arrow's theorem is. The philosophical moral is a variant of Section 5.3's: where the hairy ball theorem shows global orientation is impossible without exception, Borsuk–Ulam shows that global comparison across the whole sphere is impossible without redundancy — two structurally distinct ways the sphere's topology refuses to let a continuous global assignment behave as if the sphere were simply a very large flat sheet.

5.10 Brouwer's Fixed Point Theorem: A Third Cousin, and Its Route to Game Theory

Sections 5.8 and 5.9 placed the hairy ball theorem beside two structural cousins from other fields, Arrow's theorem and the Borsuk–Ulam theorem, in each case naming a shared argumentative shape while insisting the mathematical content itself did not transfer. A third cousin is closer to home than either, sharing not only the shape but a documented family relation in the technical literature: Brouwer's fixed point theorem, which states that every continuous function from a closed disk to itself has at least one fixed point, a point mapped to itself. The hairy ball theorem is, in the precise sense the companion volume's proof makes explicit, a close relative of Brouwer's result rather than merely an analogue of it: both follow from the same underlying machinery of degree theory, the count, with sign, of how many times a continuous map wraps one space around another, and a standard route to proving the hairy ball theorem is to show that a

nowhere-vanishing tangent field on the sphere could be used to construct a fixed-point-free self-map of the disk, contradicting Brouwer directly.

The family relation matters here for a reason beyond tidiness, because it supplies a bridge back to Section 5.8's excursion into social choice that Arrow's theorem alone could not. John Nash's 1950 proof that every finite game has at least one equilibrium in mixed strategies is, itself, an application of a generalization of Brouwer's theorem (Kakutani's fixed point theorem, for set-valued maps) to the space of strategy profiles; existence of equilibrium in game theory and the impossibility of a hairy sphere are, on this accounting, not merely similar in form but genealogically related results of the same fixed-point machinery, applied to different spaces for different purposes. Where Section 5.8 could only say that Arrow's theorem shares a shape with Chapter 5's result, this section can say something stronger and more exact about Nash's: the same mathematics that forbids a hairy sphere guarantees an equilibrium exists, the sign of the underlying degree argument reversed from an impossibility into an existence proof.

The philosophical moral is narrower than the mathematical connection might suggest, and stating the narrowness precisely is again exactly the discipline Section 5.3 asked this chapter to maintain throughout. Fixed-point theorems, as a family, are indifferent to what is held fixed — a point of a disk, a strategy profile, a vector at a point of the sphere — and prove existence or, in the hairy ball's special topological setting, nonexistence, from the shape of the space alone. It licenses no inference from Chapter 5's cartographic impossibility to any claim about economic equilibrium beyond the bare fact that both are theorems of the same mathematical family; what it does license, and what this section has tried only to establish, is that the recurring pattern Part II keeps finding — an impossibility or an existence result forced by the shape of a space rather than negotiated by anyone's choice — is not confined to topology narrowly conceived, but names a piece of mathematics with genuine reach across fields that do not otherwise resemble one another.

5.11 A Meteorological Instance: Why Global Wind Must Vanish Somewhere

Section 5.7 already moved the hairy ball obstruction out of metaphor and into a physical instance, nematic liquid crystals forced by the same parity argument to carry topological defects on a spherical droplet. A second physical instance belongs beside it, closer to ordinary experience than a liquid crystal and no less rigorous: the horizontal wind field at the Earth's surface, treated at any single instant as a continuous tangential vector field on a sphere, is subject to exactly the theorem Section 5.2 states for cartography. At every moment there must exist at least one point on the globe, and by the parity argument of Section 5.5 in general an even number of them, at which the surface wind velocity is exactly zero, a point of dead calm forced not by any particular meteorological configuration but by the topology of the sphere itself. Idealizations are doing real work in this claim, a wind field is not perfectly continuous at every scale, and 'surface wind' is itself a modeling choice about which layer of the atmosphere counts, and Section 5.3's discipline about keeping metaphor and argument apart applies here as much as anywhere else in the chapter: the theorem constrains the idealized continuous field meteorologists build from observations, not the observations themselves, which are finite, noisy, and never quite continuous.

With that qualification in place, the instance is a genuine, textbook-level illustration rather than a borrowed figure of speech, and meteorologists have long known it under a more practical name than Section 5.1's abstract statement: a cyclone's eye, or an anticyclone's center, is not merely a region of low or high pressure but a point, or a small closed region collapsing toward one, at which the wind's horizontal component vanishes, exactly where a comb drawn across a hairy sphere must leave a cowlick. This does not explain why cyclones form where and when they do, a question belonging to fluid dynamics and thermodynamics far outside this volume's scope, but it does explain something the theorem alone predicts and meteorology independently confirms: a planet with a continuously varying surface wind field can never be uniformly windy everywhere at once, no matter how vigorous the general circulation, for precisely

the reason Section 5.2 gives for the impossibility of a nowhere-zero field of fixed directions on a two-sphere. The same obstruction that forbids a universally consistent compass bearing toward Mecca, Chapter 7 will show, forbids a planet-wide gale with no calm center anywhere upon it.

Chapter 6 — Geometric Conventionalism: Poincaré and the Not-Quite-Arbitrary Choice of Projection

6.1 A question Chapters 4 and 5 leave open

Chapters 4 and 5 established what is impossible and what is merely costly, but neither chapter says how a cartographer is supposed to choose among the options that remain once impossibility has ruled some out. A navigator needs conformality; a demographer needs equal area; the theorems are silent on which of these needs is more legitimate, because legitimacy of purpose is not a geometric property. This chapter asks whether the resulting freedom — real freedom, since the mathematics alone does not determine the choice — is therefore arbitrary in the more troubling sense of being groundless, a matter of taste with no rational constraint at all. The answer this chapter defends, following a structure of argument owed to Henri Poincaré, is no: the choice is underdetermined by geometry without being ungoverned by reason.

6.2 Poincaré's conventionalism, narrowly stated

Poincaré's argument, developed in the context of the choice between Euclidean and non-Euclidean geometries as descriptions of physical space, holds that no possible experiment can force that choice on its own: any experimental result can be accommodated either by adopting a non-Euclidean geometry with simple physical laws, or by keeping Euclidean geometry and complicating the physical laws (introducing universal forces that deform measuring instruments in just the way needed to save appearances). The choice between the two packages is therefore a convention, not a discovery — but Poincaré is explicit that

'convention' does not mean 'arbitrary in the sense of unmotivated.' Euclidean geometry is, in his account, preferred because it is simpler, and simplicity is a genuine, if non-empirical, ground for choice among empirically equivalent alternatives.

The cartographic case is structurally close to this without being identical to it, and the difference is worth marking precisely. Poincaré's two packages (geometry plus physics) are empirically equivalent — no observation distinguishes them. Two map projections are not empirically equivalent in that sense: a Mercator chart and a Gall–Peters chart make different, checkable claims about angle and area, and Chapters 4 and 5 supply the exact theorems by which those claims can be verified. What is underdetermined is not which projection is correct — none is, uniquely, since none can be both conformal and equal-area — but which correctness is wanted. The conventionality here attaches not to the mathematics, which is fully determinate, but to the selection of a criterion (angle-preservation, area-preservation, some weighted compromise) by which to judge success. This is Poincaré's structure transposed: not empirical equivalence masking a choice of description, but geometric exhaustiveness (the trade-off is fully classified) masking a choice of purpose.

6.3 Why the choice is constrained without being forced

If the choice of projection were simply a matter of stated purpose, it would be trivially non-arbitrary in a thin sense (given the purpose, the projection follows) while remaining arbitrary in a thicker sense (nothing constrains which purpose is legitimate to hold). The more interesting constraint, and the one this chapter wants to isolate, is that purposes themselves are not all on an equal footing once the theorems of Chapters 4 and 5 are known. A navigator who does not know that no chart can be both conformal and equal-area might select a projection by habit or authority, unaware that a real trade-off is being made; a navigator who does know is choosing with full information about exactly what is being given up and why it must be given up. The theorems do not select the purpose, but they discipline the choice by making its cost fully explicit and, crucially, by

ruling out certain purposes as simply incoherent — 'a chart that is both conformal and equal-area' is not a demanding purpose but an impossible one, and knowing the theorem forecloses it before any design work begins. Rational constraint on an underdetermined choice, in other words, need not take the form of narrowing the choice to one option; it can take the form of making every remaining option's price fully legible, and eliminating those that were never really options at all.

6.4 From cartography back to convention in general

The suggestion this chapter offers, again cautiously, is that this is the general shape a defensible conventionalism should take wherever it is invoked: not 'the matter is undetermined by the facts, therefore any choice is as good as any other,' but 'the matter is undetermined by the facts in a specific, boundable way, and the theorems that establish the underdetermination simultaneously specify exactly what is and is not still at stake in the remaining choice.' A conventionalism that stops at underdetermination, without doing the further work of showing what is and is not thereby settled, earns the charge of arbitrariness it is usually raised to answer. Poincaré's own argument does this further work for the geometry-versus-physics case; the cartographic case, thanks to the full classification available from Chapters 4 and 5, allows the same further work to be done with unusual completeness, because there the space of remaining options is not merely known to be nonempty but exactly enumerated.

6.5 Underdetermination Without Relativism

Section 6.3 argued that Poincare's conventionalism constrains the choice of geometry without forcing it, and left the sense of 'constrained' largely to do its work by contrast with an unstated alternative, full relativism, on which the choice would be answerable to nothing beyond preference. It is worth making that contrast explicit, because conventionalism is one of the most frequently misread positions in the philosophy of science precisely along this axis, collapsed by critics into a relativism it was designed to avoid and claimed by some of its less careful successors as license for exactly that relativism. Poincare's own argument turns on empirical underdetermination: a given body of measurements is compatible with more than one combination of a geometry and a physics, because any apparent geometric anomaly can always be absorbed into a compensating adjustment to the physical laws describing the measuring instruments themselves, a rigid rod that appears to behave non-Euclidean can equally be described as a Euclidean rod subject to a universal force affecting all measurement alike. This underdetermination is genuine and, on Poincare's own analysis, permanent: no future experiment resolves it, because any experiment's result can always be redistributed between geometry and physics in the same way. But underdetermination among finitely many, structurally specific alternatives is not the same claim as the relativist's stronger one, that no alternative is answerable to anything at all. Convenience, simplicity, and the accumulated economy of an entire physical theory built on one choice rather than another remain fully operative reasons, on Poincare's account, for preferring Euclidean geometry in ordinary physics, reasons that are not truth-tracking, in the sense of picking out which geometry the world 'really' has, because the argument's whole point is that there is no fact of that kind to track, but that are also not arbitrary, in the sense of being interchangeable with any other reason whatsoever.

This is exactly the shape Section 6.4 needs when it turns from cartography back to convention in general, and exactly the shape Chapter 13 will need again, more urgently, when it turns to the choice among projections: a constrained space of

live options, underdetermined by the facts that fix the space itself, within which some further consideration, convenience, purpose, historical accident, or interest, does the actual choosing, and can therefore always be asked to justify itself, without that demand ever threatening to collapse the underlying constraint into nothing at all.

6.6 Reichenbach's Coordinative Definitions: A Second Conventionalism, Independently Stated

Section 6.2 built this chapter's conventionalism from Poincare's own argument for the choice between Euclidean and non-Euclidean geometry; a second and closely related figure in the same tradition, Hans Reichenbach, gives the position a sharper technical vocabulary worth setting alongside Poincare's, because Reichenbach's version isolates exactly the feature of the cartographic case Section 6.2 needed but did not fully name. Reichenbach's Philosophy of Space and Time distinguishes physical geometry from a coordinative definition: before any measurement can determine which geometry a physical space has, some convention must first stipulate what will count as a rigid measuring rod, a stipulation that cannot itself be settled by measurement without circularity, since the very instrument being calibrated is the one doing the measuring. Once a coordinative definition is adopted — rigid rods are ones unaffected by their motion through the field, for instance — the resulting geometry becomes a genuine empirical question; but the coordinative definition itself is chosen, not discovered, and different coordinative definitions can redescribe the same body of observations using different geometries, exactly Poincare's underdetermination restated with the specific mechanism, the choice of a coordinating convention, made explicit rather than left implicit in the notion of a 'convention' generally.

The cartographic case supplies its own coordinative definition, and naming it sharpens Section 6.3's claim that the choice among projections is constrained without being forced. Before any projection can be judged more or less faithful, a coordinative choice already fixes what 'faithful' will mean for the judgment that

follows — angle-preservation, or area-preservation, or some stated weighting of the two — and this choice, like Reichenbach's choice of a rigid rod, is not itself determined by the geometry it goes on to measure success against; Tissot's theorem presupposes, rather than settles, which of a and b (Section 4.1) a given map is trying to keep close to one. This is Reichenbach's structure transposed a second time, alongside Poincare's in Section 6.2: not a claim that cartographic success is arbitrary, but a claim that the standard of success is itself a coordinative definition, chosen prior to and independent of the theorem that then rigorously governs how well any given projection meets it — the same two-stage structure, convention first, rigor second, that Reichenbach isolated for the choice of geometry itself.

6.7 Convention as Coordination: Lewis's Different Sense of the Word

Sections 6.2 through 6.6 have used 'convention' in the epistemic sense Poincaré and Reichenbach gave it: a stipulation, such as which quantity a projection is asked to preserve, that is not fixed by the geometry it goes on to be judged against, though the judgment itself, once the stipulation is made, is fully rigorous. A different and independently developed sense of the same word is worth setting alongside it before Part II closes, because the two are easy to run together and answer different questions. David Lewis's Convention: A Philosophical Study analyzes convention game-theoretically, as a regularity in the behavior of a population, sustained because each member expects the others to conform and prefers to conform given that expectation, in a situation admitting more than one such self-sustaining regularity — a coordination problem with multiple equilibria, any one of which would serve the population equally well once generally adopted. Which side of the road to drive on is Lewis's standard illustration: nothing about the physics of driving favors the right over the left, but once nearly everyone drives on one side, everyone has a standing reason to continue, and the reason has nothing to do with rigor or measurement at all.

Cartography supplies conventions of both kinds, and distinguishing them prevents a specific confusion the chapter has so far had no occasion to name. The choice of what a projection preserves, angle or area or some declared compromise (Section 6.3), is a Poincaré-Reichenbach convention: underdetermined by the mathematics, but once made, subject to Tissot's theorem without further slack, a rigorous consequence with no equilibrium structure to it at all. That north sits at the top of nearly every printed map, by contrast, is a Lewis convention in something close to the textbook sense: no projection requires it, no theorem constrains it, a map with south at the top is neither less conformal nor less equal-area for the reversal, and the practice persists because readers expect it and mapmakers have every reason to meet the expectation, exactly the self-sustaining structure Lewis's account describes — with genuine historical counter-equilibria to prove the point, since south-up orientation was the norm in a number of medieval Islamic and East Asian cartographic traditions, oriented instead toward Mecca or toward the emperor's seat, a different but equally stable equilibrium serving a different community's coordination needs.

The two senses are not competitors for the same territory; each governs a different part of what a map is. Reichenbach's coordinative definition fixes the standard a theorem then holds the map to; Lewis's coordination equilibrium fixes a practice no theorem touches at all. Calling both 'convention' risks suggesting that the north-up practice is, like the choice between conformal and equal-area, a stipulation later vindicated or constrained by geometric proof; it is not, and no proof is available to vindicate it, because there is nothing for a proof to be about. Keeping the two apart is the same discipline the Preliminary Note asked for at the level of a single word: not every convention is a coordinative definition, and Part II's mathematics has something rigorous to say about only one of the two kinds.

6.8 Carnap's Principle of Tolerance: Conventionalism Generalized Beyond Geometry

Sections 6.2 through 6.8 have kept Poincaré's conventionalism narrowly enough stated that it applies to the choice of geometric description without licensing the stronger claim that any description is as good as any other; Reichenbach's coordinative definitions, Lewis's convention as coordination, and Grünbaum's metrical amorphousness were each shown to occupy a different, carefully bounded position along that spectrum. Rudolf Carnap's principle of tolerance, stated in Logical Syntax of Language, generalizes the conventionalist move one level further, from geometry to language itself: for Carnap, the choice among logical syntaxes, among what he calls linguistic frameworks, is not a question to be settled by argument about which framework is uniquely correct, but a pragmatic decision each investigator is free to make, provided only that the chosen syntax is stated with full explicitness and applied consistently thereafter — in his own later formulation, everyone is at liberty to build their own logic, their own form of language, as they wish.

The cartographic case supplies Carnap's principle an instance narrower and more tractable than the language cases he actually discussed, and for that reason a useful test of how far the generalization travels. Choosing a map projection is, in Carnap's vocabulary, an external question, a choice of framework, answerable in terms of the framework's convenience, fruitfulness, and fit to purpose rather than in terms of which projection is uniquely true of the earth; asking whether a given distance on the Mercator projection is accurate is, by contrast, an internal question, answerable within the chosen framework by ordinary geometric computation once the projection is fixed. Section 6.4's move from cartographic convention to convention in general already anticipated something like this distinction without naming it; Carnap's internal/external pairing supplies the missing general vocabulary for exactly the difference Section 4.3 drew, in cartography's own terms, between disputing which projection to adopt and disputing whether a computation is correct within an already-adopted one.

The principle of tolerance is not, for all that, a license for the relativism Section 6.5 was careful to keep conventionalism apart from. Carnap's tolerance operates only at the level of framework choice; once a framework, a projection, a

coordinate system, a logical syntax, is adopted, every internal question within it has a determinate answer, and Tissot's theorem is exactly such an internal, determinate consequence once the framework of continuous projection from sphere to plane is fixed. What Carnap adds to this chapter's argument is the observation that the same two-level structure, freedom at the level of framework, constraint within it, that Chapters 4 through 6 established for cartography specifically is not a peculiarity of maps; it is, on Carnap's account, the general shape any formal representational choice takes, which is one further reason cartography has served this volume as well as it has as a test case for representation beyond itself.

6.9 Quantum No-Cloning as a Third Register of Representational Impossibility

Sections 4.1–4.6 and Chapter 5 established two impossibility results in this volume's core domain: a forced trade-off among incompatible virtues, and a topological obstruction to a specific kind of global consistency. A third and physically more exotic register, useful for testing how far the pattern generalizes, is the no-cloning theorem of quantum mechanics: no physical process can produce an exact copy of an arbitrary, unknown quantum state, a direct consequence of the linearity of quantum evolution, proved independently by Wootters and Zurek and by Dieks in 1982.

No-cloning sits in a different position from either earlier case. Tissot's theorem (Chapter 4) permits a genuine choice among incompatible virtues once one accepts that all cannot be had; the hairy ball theorem (Chapter 5) forbids a specific global structure but says nothing against locally consistent alternatives, several of which Chapter 7 will examine. No-cloning forbids something prior to any of this: not a particular representation of a state, but the very possibility of obtaining a second copy to represent, compare, or select among in the first place. And yet a conventional residue survives even here, in the sense Section 6.2 gives Poincaré's conventionalism: which basis one measures an unknown state in remains a free choice, even though the state cannot be duplicated for

measurement in more than one basis — constraint without full determination, at one further remove.

The discipline of the Preliminary Note applies with particular force to this case, since quantum mechanics has attracted more loose metaphor per genuine result than almost any other source this volume could draw on. No-cloning is a fact about physical systems, not about pictures of physical systems, and its relevance here is exclusively structural: a third, independently provable instance, after Tissot's trade-off and the hairy ball's topological block, of a representational impossibility that is a theorem rather than a truism — offered only as that, and nothing broader.

6.10 Duhem–Quine Underdetermination and the Projection Choice

Section 6.2 read Poincaré's conventionalism narrowly, as holding that geometric facts fix a trade-off within which some choice remains free. A related but distinct underdetermination thesis, originating with Pierre Duhem's 1906 study of physical theory and extended by Quine into a general holism, holds that no theory is tested in isolation: any observation that conflicts with a theory's predictions can, in principle, be accommodated by revising some other auxiliary assumption in the surrounding web of belief instead of abandoning the theory itself, so that evidence alone never uniquely selects which piece to revise.

The two theses share a family resemblance worth stating precisely rather than assuming. Duhem's own 1906 discussion, in fact, already treated the choice of geometry for physical space as a paradigm case of exactly this holist underdetermination, making him a direct ancestor of the position Section 6.6 credits to Reichenbach's coordinative definitions. But the structures differ where it matters: this volume's cartographic conventionalism (Sections 6.2–6.8) concerns freedom among options that remain equally compatible with a single, already-settled body of geometric fact — Tissot's theorem has already told us exactly what any projection can and cannot achieve, and convention operates only inside that closed mathematical result. Duhem's and Quine's

underdetermination instead concerns which empirical hypothesis to revise when the facts themselves remain genuinely open to more than one accommodation.

What is gained by placing cartographic conventionalism inside this larger family is a sharper sense of how narrow, and how unusually well-behaved, the cartographic case is. It is conventionalism with the empirical question already closed — curvature is not in doubt, and no future observation will revise Tissot's theorem — paired with a genuine, bounded freedom in how to depict a settled fact. Quine's general holism, by contrast, usually operates precisely where the empirical question remains open, which is why underdetermination in his sense is felt as a live epistemological difficulty in a way that Section 6.5 already argued cartographic underdetermination need not be.

6.11 A Case in Point: The EPSG Registry and Convention Made Infrastructure

Section 6.7 distinguished Lewis's sense of convention, a solution to a coordination problem stable only because everyone expects everyone else to keep using it, from Poincaré's sense, a genuinely free choice among empirically equivalent geometric descriptions. A single, well-documented technical registry shows what Lewis's sense looks like once it has to be implemented in software rather than merely observed in a population's behavior. In 1985 a working group within the European Petroleum Survey Group, an association of oil companies needing to exchange survey and seismic data across national boundaries and internal projection conventions, began assembling what became the EPSG Geodetic Parameter Dataset, a numbered catalogue in which every geodetic datum, ellipsoid, coordinate reference system, and transformation between systems is assigned a unique integer code. The dataset was made public in the early 1990s, with an initial release of a few hundred definitions, and in 2005 the European Petroleum Survey Group itself was dissolved and its maintenance responsibilities absorbed into the International Association of Oil & Gas Producers' Geomatics Committee, which kept the original acronym rather than rename a registry already in wide use. The code most readers of this volume will

have already encountered without knowing its origin, EPSG:4326 for the latitude-longitude system underlying the Global Positioning System, is one entry among many thousands, each a small, arbitrary-looking integer standing in for a full technical specification no two pieces of software could otherwise be trusted to agree on.

What the registry adds to Section 6.7's account is not a new philosophical point but a demonstration of an old one under load: nothing internal to geodesy forces EPSG:4326 to be four digits beginning with a four rather than any other numeral, and nothing about oil exploration bears on why a global navigation system should care about a mapping standard first assembled to reconcile seismic surveys, and yet, once thousands of software libraries, government agencies, and web-mapping platforms had built EPSG codes into their own systems, changing the numbering would carry the same practical cost Section 6.7 attributes to any Lewis-convention once it is entrenched, an ordinary technical necessity, not a discovery about the shape of the Earth. This is the coordination problem Chapter 4's compromise-projection choices and Chapter 6's own conventionalism have both described in the abstract, made visible as an actual piece of infrastructure: a numbered Schelling point born inside a petroleum consortium's internal bookkeeping and now silently underwriting a large fraction of the world's digital cartography, constrained, as Section 6.3 put it of projection choice generally, without being forced, and no less binding in practice for that.

PART III

A Phenomenology of Orientation

Chapter 7 — The Retroazimuthal Problem as a Structure of Dwelling

7.1 From the third person to the first person

Chapter 5 stated the hairy ball obstruction impersonally: a continuous field of directions toward a fixed point C cannot be defined on the whole sphere. Stated this way, the result belongs to no one in particular; it is a fact about a geometric object. But the situation it describes is, in every historically important instance, someone's situation: a person standing at a particular point P, wanting to know, right now, which way to face. This chapter asks what changes when the third-person theorem is put back into the first-person circumstance that motivated its cartographic study in the first place — not to weaken the theorem into anecdote, but to notice that the theorem was never, historically, pursued for its own sake. It was pursued because people needed an answer.

The companion volume The Retroazimuthal Problem treats this need in its most sustained historical form: the medieval Islamic science of determining the qibla, the direction of the Kaaba in Mecca, from any locality, for the purpose of ritual prayer. That volume's concern is mathematical and historical — the spherical trigonometry of the problem (Chapter 2 there), the astronomers who solved it with increasing sophistication (Chapter 4), and the three cartographic constructions, centuries later, that solve the general problem of displaying the answer on a map (Chapters 8 through 11 there). This chapter's concern is different: not how the qibla was computed, but what kind of relation to space the need to compute it expresses, and whether that relation is peculiar to ritual prayer or a more general structure of what it is to dwell somewhere at all.

7.2 Facing, not merely locating

An ordinary act of locating oneself — reading a coordinate, finding a position on a map — answers the question 'where am I.' The retroazimuthal need answers a different question: 'which way must I turn, from here, to face something I am not

at.' The difference is not merely grammatical. Locating oneself is satisfied by a single fact, one's own coordinates; facing a fixed center requires, in addition, a relation between that single fact and every other point on the sphere, because the correct answer to 'which way' changes continuously as P moves, in a manner governed by the asymmetry of spherical bearing discussed at length in the companion volume (its Chapter 2, on the signed spherical excess and the non-reversibility of bearing under exchange of the two points). To face a center is thus not a local fact about where one stands but a relational fact about where one stands with respect to everywhere else one might have stood instead — which is exactly the structure that forces the hairy ball obstruction to bite: a coherent, everywhere-consistent assignment of 'which way to face C' is precisely a tangent vector field on the sphere, and Chapter 5 has already shown why no such field can be smooth and non-vanishing everywhere.

This suggests a philosophical redescription of the qibla problem that does not depend on any specifically religious content: it is the general problem of maintaining a stable orientation toward an elsewhere, from wherever one happens to be, without that elsewhere ever becoming one's own location. The center is defined precisely by not being occupied by the one orienting toward it. A pilgrim en route to Mecca, on arrival, no longer has a qibla to determine — the question dissolves exactly at the center, which is the geometric content of the singularity the theorem forces there. Facing is a structure available only to the not-yet-arrived.

7.3 Two singularities, two kinds of impossibility

Chapter 5 noted, following the companion volume, that the retroazimuthal field is forced to fail not at one point but at two: the center C itself, where facing collapses because there is no longer a direction to determine, and the antipode of C, where every direction is equally, and therefore no longer meaningfully, a direction toward C. These two singularities are not philosophically equivalent, though the theorem treats them symmetrically. Arrival at the center is the resolution of the orienting project — the question stops applying because it has

been answered by presence. Arrival at the antipode is a different kind of failure: the question still seems to apply, one still wants to know which way to face, but the sphere's geometry no longer supplies a fact of the matter, because every direction leaving the antipode approaches C by a great-circle path of exactly the same length and, in the relevant limit, the same claim to being 'the' direction. The center's singularity is a completion; the antipode's is a genuine loss of orientation, the one point on the sphere from which facing C is not merely difficult but strictly undefined. Historically, this is also why the far antipodal region from Mecca (in the Pacific) played essentially no role in the medieval qibla literature the companion volume surveys — not because no one lived there, but because the question that literature answers has no answer there to find.

7.4 What Part III will do with this

The remainder of Part III keeps this dwelling-structure in view while adding two further cases that extend it in different directions. Chapter 8 stays with the qibla itself, but turns from the geometry of facing to the specific historical episode in which that geometry was worked out — the medieval astronomers' own methods, and what the choice between an exact and an approximate method reveals about the relation between ritual obligation and mathematical rigor. Chapter 9 leaves fixed orientation toward a point for a different, but structurally related, first-person problem: the stability of the ground one stands on, taken up through the recent modernization of the geodetic datum (NATRF2022, discussed technically in the companion volume The New Datum) as a case of what it means for the reference frame itself — not one's bearing within it — to be found in motion.

7.5 The Traveler's Version: Dead Reckoning and the Ordinary Phenomenology of Being Lost

This chapter's argument has been built, so far, on the single most historically developed instance of the retroazimuthal need, the qibla; it is worth showing briefly that the structure identified there is not specific to ritual orientation, by

examining its secular twin. A ship's navigator using dead reckoning, estimating present position from a known starting point, a heading, a speed, and an elapsed time, without external fixes, is solving a cousin of exactly the problem Section 7.2 described, and the two singularities of Section 7.3 have direct, if less dramatic, analogues in the navigator's practice. There is no religious center C in dead reckoning, but there is a destination, and the entire discipline is organized, like the qibla problem, around maintaining a relation to a point one does not currently occupy; the practice's traditional name for the moment that relation is lost, when accumulated error in heading and speed estimates has grown large enough that the navigator no longer has confidence in the ship's position relative to any charted point, is being 'out of one's reckoning', a phrase whose ordinary-language descendant, simply 'being lost', still carries the same structure Section 7.2 formalized: not a failure to know where one is, in isolation, but a failure of the relation between where one is and where one is going.

What the secular case adds to the religious one is a clearer view of the antipodal singularity's ordinary-language cousin, because open-ocean dead reckoning has its own version of a region where the method characteristically breaks down hardest: far from any landmark, over long elapsed time, where the compounding of small heading and speed errors, the practical analogue of the geometric fact that every direction leaving the antipode is equally a direction toward C, leaves the navigator's estimate of position spread over an area large enough that no single point within it can be preferred to any other. The phenomenology of being lost in this sense, familiar to anyone who has misjudged a return route in unfamiliar terrain long before any theorem is invoked to explain it, is this chapter's dwelling-structure recognized from the inside, without the vocabulary of centers and antipodes, evidence, of the kind this volume's method treats as illustrative rather than probative, that the structure derived here from the geometry was available to ordinary experience well before geometry gave it a name.

7.6 Merleau-Ponty and the Body as the Unmarked Center

Section 7.2 described facing a fixed center as a relational fact obtaining between one's own location and every other point one might have occupied instead, and derived from this relational structure the same obstruction Chapter 5 proved abstractly. Maurice Merleau-Ponty's phenomenology of embodied space, developed at length in the Phenomenology of Perception, supplies an independent route to a closely related observation, arrived at from lived experience rather than from topology, and it is worth setting the two alongside each other with the same care Section 7.1 promised for every extension beyond the strictly geometric. Merleau-Ponty argues that spatial orientation is not first a matter of locating oneself within an already-given, homogeneous geometric space, the way a coordinate is located within a coordinate system, but that the lived body is itself the origin from which any such space becomes available at all — up, down, near, and far are not read off a neutral grid but constituted by the body's own orientation and capacity for movement within a situation, so that the body functions, in his phrase, as the 'null point' of spatial orientation rather than as one occupant of space among others. His well-known discussion of experiments with inverting goggles, in which a subject's visual field is optically flipped and the world initially appears upside-down before the subject's active engagement with it gradually restores a lived, functional 'right-side-up', is offered as evidence that orientation is achieved through embodied action rather than given passively by the geometry of the retina.

The resonance with Section 7.2 is real and worth naming precisely rather than left to gesture at itself: both accounts locate something structurally special about the point from which orientation proceeds — Merleau-Ponty's lived body as the unmarked origin of a perceptual field, this volume's point P as the unmarked, always-shifting origin of a bearing toward C — and both find that this specialness is not a contingent fact about any one body or locality but a structural feature of what it is to have an orientation at all. The disanalogy this volume's discipline requires stating is equally real: Merleau-Ponty's null point is a claim about the phenomenology of perception, established, if at all, by careful first-person and experimental description, not by proof, and admits of no

analogue to Chapter 5's theorem — nothing in the Phenomenology of Perception forces, in the strict sense Part II reserves for a genuine impossibility, that embodied orientation must fail at some point the way a tangent vector field must. What Merleau-Ponty's account offers this volume is not a second proof but a first-person description, independently arrived at by a philosopher not working with cartography at all, of exactly the asymmetry between the orienting subject and the everywhere-else it orients toward that Section 7.2 derived from the geometry — a convergence, in the modest sense Section 11.5 will later make the volume's standard for such things, not a further theorem.

7.7 Lynch's Cognitive Maps: Way-finding Without a Global Field

Section 7.5 examined dead reckoning, the sailor's practical accommodation of Chapter 5's obstruction, tracking position by accumulated estimate rather than by any continuously available global bearing; a second, more common accommodation, this one governing how ordinary people navigate ordinary cities, is worth setting alongside it before Part III's dwelling-structure closes. Kevin Lynch's The Image of the City, based on structured interviews with residents of Boston, Jersey City, and Los Angeles, found that people do not carry anything resembling a continuous, metrically accurate mental map of their city; they build and rely on what Lynch calls a cognitive map, assembled from five recurring kinds of element, paths (the streets and routes actually traveled), edges (boundaries such as a river or a rail line), districts (regions recognized as sharing some character), nodes (junctions and gathering points), and landmarks (distinctive, externally visible reference points) — a sparse, locally anchored structure, not a scaled-down Tissot-respecting projection of the city carried in the head.

The connection to this chapter's argument is closer than a loose parallel between two different senses of 'map.' Chapter 5 proved that no continuous, everywhere-consistent field of directions toward a fixed point can exist on the sphere; Lynch's empirical finding is that ordinary human navigation does not, in fact, depend on any such field, global or local, and does not seem to be missing it.

Way-finding by landmark and path, on Lynch's account, succeeds by chaining together local, overlapping fragments, each valid only in its own neighborhood and each silent about how to get anywhere outside it, precisely the kind of construction Section 5.4 anticipated Part III would need to examine: not a workaround that restores a global field by other means, but a genuinely different strategy that never required one. A resident who cannot state the compass bearing from home to the train station, and has never needed to, is not failing to solve Chapter 5's problem; the problem, posed as a request for a single global field, was never the one that way-finding, as Lynch's subjects actually practiced it, needed to solve.

This is worth distinguishing carefully from Section 7.5's dead reckoning and from Section 8.5's compass apps, because the three responses to the same underlying obstruction are not equivalent. Dead reckoning approximates a single continuous trajectory by accumulated local estimate, aiming at the same global answer the impossible field would have given, if only imperfectly; a compass app restores something close to a genuine global field computationally, for a single reference point, at the cost of the device the field is displayed on. Lynch's cognitive map does neither: it neither approximates nor computationally restores a global field, but replaces the need for one with a different representational economy entirely, built from overlapping local fragments that need never be reconciled into a single consistent whole because the tasks they serve, getting from this junction to that landmark, never require reconciliation. Of the three, it is the one that takes Chapter 5's impossibility most fully in stride, not by defeating it but by never having posed the question in the form the theorem forbids an answer to.

7.8 Wayfinding Without Instruments: The Polynesian Star Compass and Etak

Sections 7.5 and 7.7 drew their evidence from Western nautical practice and Western urban cognitive mapping respectively. A documented and independently developed solution to the same underlying problem — how to

navigate accurately across open ocean without the global field Chapter 5 rules out — comes from traditional Pacific Islander wayfinding, reconstructed and publicly demonstrated in the twentieth century through navigators such as Mau Piailug, who sailed the Hokuleʻa between Hawaiʻi and Tahiti in 1976 using no instruments at all.

The core technique, etak, inverts the usual frame: rather than tracking the canoe's motion against a fixed grid, the navigator treats the canoe as stationary and a chosen reference island, invisible below the horizon, as moving backward past it, position being read off the changing bearing to that reference rather than off any coordinate fix. This is paired with a star compass, a memorized ring of horizon positions marking where particular stars reliably rise and set, giving a stable set of directional landmarks that Section 7.7's Lynchian vocabulary would recognize as nodes, though celestial rather than urban ones.

Etak never attempts to instantiate the single continuous field Chapter 5 proves cannot exist; like dead reckoning (Section 7.5) it accumulates a running local estimate, and like cognitive mapping (Section 7.7) it is built from a chain of landmarks rather than a metric grid, but it is neither of those two things exactly — a third solution-shape, arrived at by a tradition with no historical contact with either the Western nautical or the Western urban-planning literatures this chapter has otherwise relied on, which strengthens rather than merely illustrates the claim that Chapter 5's obstruction has no single required workaround.

7.9 Heidegger's Zuhandenheit and the Map That Disappears in Use

Section 7.6 read Merleau-Ponty on the lived body as the unmarked center from which orientation proceeds. A related but distinct phenomenological resource bears on the map specifically as a piece of equipment: Heidegger's distinction, in Being and Time, between the ready-to-hand (zuhanden), a tool absorbed transparently into skilled activity, and the present-at-hand (vorhanden), the same object made an explicit object of scrutiny, typically because it has stopped working smoothly.

A well-used map behaves like Heidegger's hammer: while wayfinding goes well, the hiker glances and walks, folding the map's information into ongoing action without pausing to read it as a representation at all, in a mode Heidegger holds is not merely inattentive but the map's proper and most complete mode of functioning. It becomes present-at-hand — an object one now looks at, rather than looks through — exactly at the moment of breakdown: a landmark absent, a symbol misread, a bearing that does not match the terrain. This is also, not incidentally, the moment at which Chapters 4 through 6's technical apparatus is actually consulted by an ordinary user, who has no occasion to think about Tissot's theorem until the map has already, in Heidegger's sense, failed.

The two phenomenological resources are complementary rather than interchangeable, and the volume keeps them apart accordingly, following the discipline Section 7.3 already set for a related pair of distinctions. Merleau-Ponty locates the zero-point of orientation in the lived body itself, prior to any map being consulted at all; Heidegger's contribution concerns the map's own mode of givenness as a tool, and says nothing by itself about the body. A account of wayfinding that used only one would be missing a real and separable half of the phenomenon.

7.10 Echolocation and Blind Wayfinding: Orientation Without a Visual Field

Sections 7.5, 7.7, and 7.8 examined sighted navigation strategies — dead reckoning, Lynch's cognitive maps, and the Polynesian star compass — each differing in method but each still organized around visually perceived landmarks or bearings. A further case dispenses with vision altogether: blind and visually impaired wayfinding, including human echolocation, the deliberate use of tongue clicks and the interpretation of their returning echoes to build spatial awareness of surroundings, taught systematically by organizations built around practitioners such as Daniel Kish, himself blind from early childhood.

Echolocation builds a structure recognizably similar to Section 7.7's cognitive map — paths, edges, and landmark-like reflective surfaces — but assembled through a different sensory channel entirely, and it exhibits the same restraint Section 7.8 noted in etak: it never requires, approximates, or gestures toward the single continuous global field Chapter 5 proves impossible. Orientation is built instead from a succession of local echo-soundscapes, each triangulated relative to the navigator's own body and immediate motion, in a manner that extends Section 7.6's Merleau-Pontian point about the body as the zero-degree of orientation to a modality Merleau-Ponty's own largely visual examples did not emphasize.

This closes out a count worth making explicit, since it has been assembled piecemeal across this Part. Sighted dead reckoning, sighted cognitive mapping, instrument-free celestial navigation, tool-mediated navigation, and now non-visual echolocation stand as five independently documented solution-shapes to the single obstruction Chapter 5 proves. None of the five approximates the impossible global field more closely than the others; each is simply a different, locally sufficient answer to a question the sphere's topology forbids from having any single, uniform answer at all.

7.11 A Contemporary Coda: Turn-by-Turn Navigation and the Skill Heidegger's Hammer Still Required

Section 7.9 read Heidegger's Zuhandenheit into the paper map used skillfully: a competent walker consulting a map to reach an unfamiliar destination is not, in the moment of fluent use, attending to the map as an object at all, the map recedes into the transparent background of the activity the way a carpenter's hammer recedes for a carpenter mid-stroke, present-at-hand only when it fails, folds wrong, or shows a road that has since closed. What that reading depended on, without quite stating it as a separate claim, is that the skill the map's disappearance presupposes, relating a two-dimensional symbolic representation to a three-dimensional terrain, holding a heading in mind, updating position by dead reckoning between landmarks, remains fully intact and simply idle while

the map is working smoothly; the walker could, if the map failed, fall back on exactly the capacities Section 7.5 described dead reckoning as requiring. Turn-by-turn satellite navigation, now the dominant mode by which most people in wealthy countries travel unfamiliar routes, achieves a superficially similar disappearance, the phone or dashboard display recedes from attention just as the paper map did, but by a different and philosophically opposed mechanism: it does not let a skilled relation to a representation become transparent through practice, it replaces the representation-relating task itself with a sequence of discrete imperative instructions, turn left in two hundred meters, that require no cognitive map, no heading, and no dead-reckoning capacity to execute correctly.

The difference is not merely armchair phenomenology. A 2020 study from McGill University tracking fifty regular drivers found that greater lifetime use of GPS navigation was associated with measurably worse performance on self-guided, unaided wayfinding tasks, and a three-year longitudinal follow-up on a subset of the same participants found that increased GPS use over that interval predicted a steeper decline in the specific, hippocampus-dependent spatial memory strategy Section 7.6's account of the body as unmarked center presupposes; a companion commentary in Nature the same period urged, in its own title, that habitual users of such systems risked losing rather than merely resting the underlying navigational skill. This is the retroazimuthal structure of dwelling in reverse: where Sections 7.5 through 7.10 traced skills, the traveler's dead reckoning, the Polynesian navigator's etak, the blind wayfinder's echolocation, that intensify a first-person orientation capacity precisely because no external representation is available to offload it onto, turn-by-turn navigation offloads the capacity so completely that Heidegger's hammer stops being a useful figure for what happens when the device fails; a walker whose phone dies mid-route, unlike Section 7.9's walker whose map blows away, often has no dormant skill left to fall back on, only the bare fact of being lost with no method for becoming otherwise.

Chapter 8 — Sacred Orientation and Cartographic Orientation: The Qibla as a Case Study

8.1 A distinction the geometry alone cannot make

Chapter 7 treated the qibla as a single well-posed geometric question: given a locality P, find the direction toward the fixed center C. But this description already smooths over a distinction the historical record insists on keeping open: the difference between finding the direction and finding it exactly. The hairy ball theorem, as invoked in Chapter 5, is a statement about the impossibility of a globally continuous exact assignment; it says nothing about how close an approximate, locally defined assignment must come before it counts, for the purpose at hand, as having answered the question. That further question — how much error a direction can carry and still be the right direction — is not answered by the geometry at all. It is answered, historically, by two different communities working on the same problem with two different standards of success: the astronomers who computed the qibla, and the jurists who ruled on what a worshipper was obligated to do with the computation.

8.2 Two methods, one problem

The companion volume The Retroazimuthal Problem sets out, in its own technical register, the two families of solution the medieval sources record (its Chapters 4 and 8 through 11 give the trigonometry and the later cartographic constructions in full). What matters here is the qualitative difference between them, not the derivations. One family — the analogical or approximate methods, some no more than a rule of thumb relating the qibla angle to the difference in latitude and longitude between locality and Mecca — gives a usable direction cheaply, at the cost of an error that grows, unevenly and sometimes badly, as the locality moves away from the small set of reference cities for which the rule was calibrated. The other family, associated with the fully worked-out spherical-trigonometric solution reached by the later astronomers, gives the exact

bearing everywhere, at the cost of a calculation few outside the observatories could perform unaided. Both methods solve, in a narrow technical sense, the same problem stated in Chapter 7: which way to face a fixed C. They do not, however, offer the same thing to the person who has to act on the answer.

8.3 What jurisprudence did with the astronomers' exactness

It would be a natural but mistaken assumption that the exact method, once available, simply displaced the approximate ones as a better calculator displaces an abacus. The jurisprudential literature took the opposite path. It is a documented and much-discussed feature of the classical legal schools that the majority position does not require the worshipper to face the exact bearing (what the sources call 'ayn al-qibla, the eye of the qibla) but only the general direction (jihat al-qibla) — with exactness reserved, as a stricter obligation, for the immediate vicinity of the Kaaba itself, where the target is no longer a point at astronomical distance but a building one can see. This volume does not adjudicate among the schools' several formulations of how wide 'general' is allowed to be, nor does it take a position on the underlying legal reasoning; the historiography here is genuinely contested, and nothing in the geometry settles it. What is philosophically significant is only the shape of the outcome: a community in possession of an exact method did not make exactness the standard of ritual adequacy. It treated the astronomers' rigor as available rather than obligatory — a resource one could draw on, especially where the approximate methods failed conspicuously, without making its absence a defect in the prayer of someone who lacked access to it, or the mathematics to use it.

8.4 Rigor as an offering, not a requirement

The lesson this episode offers the volume's larger argument is not that ritual obligation is indifferent to truth, which would misdescribe both the astronomers' evident seriousness and the jurists' evident respect for their results. It is that a practice can hold two different relations to the same piece of exact knowledge at

once: needing it to exist, so that the correct direction is a fact and not a convention to be argued about locally, while not requiring every practitioner to possess it, so that the obligation remains performable by anyone, anywhere, with only an estimate. Mathematics, left to itself, does not grade its own results by how performable they are; a proof is exact or it is not, and Chapter 4's statement of Tissot's theorem admits no partial credit for a projection that is only slightly conformal. Ritual obligation, by contrast, appears able to accept a graded standard without thereby becoming careless about the fact that a single true bearing exists. The qibla case suggests that this is not a compromise forced on religion by human limitation, to be corrected as instruments improve, but a stable and in a sense reasonable division of labor between two different questions: what is the case, which admits of exactness, and what may be asked of a person situated as this person is, which does not have to inherit that exactness to remain a serious question. Part III's next case, the drifting datum of Chapter 9, will show a structurally opposite difficulty: a situation where even the astronomers' 'exact' answer turns out to have an expiration date built into it.

8.5 A Modern Echo: Compass Apps and the Return of Exactness

Sections 8.2 through 8.4 described a historical situation in which the exact method for finding the qibla required an astronomer's training and instruments, while the approximate methods, cheap and widely usable, carried an error the jurisprudential majority position did not require worshippers to eliminate. That situation has changed in a way worth pausing on before Chapter 9 moves to a different case, because a smartphone's GPS receiver and digital compass now compute a locality's exact bearing to the Kaaba, to a precision no medieval astronomer could have matched, at effectively zero cost to anyone who owns the device, collapsing, for a large fraction of the world's Muslim population, the very asymmetry between 'ayn al-qibla and jihat al-qibla that Section 8.3 treated as historically load-bearing. It is worth asking directly whether this technological change unsettles the philosophical point that section drew from it.

It does not, though seeing why requires distinguishing two things the historical asymmetry was doing at once. One was practical: exactness used to be expensive, so the graded standard functioned, among other things, as a workaround for a real access problem, and where the access problem has genuinely disappeared, the workaround's practical necessity disappears with it, a worshipper with a functioning phone today has little reason to prefer a rule-of-thumb over the app's exact bearing, and there is no obvious jurisprudential reason to insist on the older, worse method now that the better one costs nothing. The other thing the asymmetry was doing, though, was not practical but structural, and survives the change intact: the ruling that jihat al-qibla suffices was never framed, in the sources this volume has followed with appropriate caution, as a concession to the poverty of available instruments, to be withdrawn once instruments improved, but as a judgment about what the obligation itself requires of a person situated at a distance from the center, a judgment about the relation between ritual adequacy and exact knowledge, not merely about the cost of acquiring exact knowledge. The smartphone changes the second half of Section 8.2's comparison, the cost of the exact method, dramatically and for almost everyone; it leaves the first half, what the obligation demands, exactly where the jurisprudence already put it. That the two can move independently, one collapsing with technology while the other stays fixed by a different kind of reasoning entirely, is itself the clearest evidence available that Section 8.4's 'offering, not a requirement' was never simply a description of medieval scarcity.

8.6 A Comparative Case: Facing Jerusalem

Section 8.3 examined a single jurisprudential tradition's resolution to the exactness question the qibla poses; a structurally similar problem, addressed independently within a different legal and ritual tradition, is worth a brief comparative mention, not to build an argument on it this volume has not done the historical work to sustain, but to test, in the modest way Section 11.5 will later formalize as this volume's standard for such comparisons, whether the resolution Section 8.4 drew from the Islamic legal sources is peculiar to that

tradition or a more general response available wherever a ritual obligation is stated in directional terms. Rabbinic Judaism requires that the Amidah, the central prayer of Jewish liturgy, be recited while facing Jerusalem — and, more specifically according to later authorities, facing the site of the Temple within it — from wherever the worshipper stands, a requirement structurally identical to the qibla's: a fixed terrestrial center, a worshipper at an arbitrary point P, and an obligation to face it. The halakhic sources addressing this requirement, discussed across the Talmudic and later codificatory literature, similarly distinguish an ideal, exact orientation from what is actually required of a worshipper without access to precise means of determining it, and similarly treat the requirement as satisfied by a good-faith orientation toward the general direction rather than voided by an error that exact astronomical or geodetic calculation would later reveal.

This volume does not have the specialized command of the halakhic literature that would let it press the comparison further than this general structural parallel, and it deliberately does not try to; the point being made is narrower and more defensible than a claim about the two traditions' legal reasoning, which this volume is not positioned to adjudicate. It is that two independent legal and ritual traditions, developing their directional-prayer requirements without reference to each other and addressing two different fixed centers, arrived at the same shape of resolution Section 8.4 described for the Islamic case: an obligation that presupposes a single fact of the matter about the true direction, while not making a worshipper's ignorance of that fact, absent the means to determine it exactly, a defect in the prayer. That two traditions facing the identical geometric problem — a fixed center, an obligation to face it, and unevenly distributed access to the means of computing the bearing exactly — converge independently on the same graded standard is modest evidence, in exactly the sense this volume has tried to earn such evidence throughout, that the resolution described in Section 8.4 is not an idiosyncrasy of one legal tradition's relationship to mathematics, but a stable response available wherever the underlying geometric structure Chapter 7 described is itself present.

8.7 Eliade's Axis Mundi: A Sacred Center That Is Ontological, Not Only Geometric

Section 8.6 compared the qibla to the Jewish practice of facing Jerusalem in prayer, holding the two cases to the same geometric structure Chapter 7 established while noting the differences in historical method. A further comparison, drawn from the general history of religions rather than from any single tradition's jurisprudence, sharpens what is distinctive about that geometric structure by contrasting it with a center of a different kind. Mircea Eliade's account of sacred space, most fully worked out in The Sacred and the Profane, describes religious experience as constituting space around one or more axis mundi, a sacred center, a temple, a mountain, a world-tree, that is not merely a point among other points but the place where, for the community that recognizes it, communication between cosmic planes becomes possible at all; profane space, on Eliade's account, is homogeneous and directionless in itself, and it is the irruption of the sacred at a center that first breaks that homogeneity and gives the surrounding space an orientation, a real up and down, a real here relative to an absolute elsewhere.

The comparison with Chapter 7's geometric center is instructive precisely where it fails to line up cleanly, and the failure is worth stating rather than smoothed over. The qibla's center, the Kaaba, is a single, geodetically locatable point, and the obstruction Chapter 5 proves against a global field of bearings toward it is a fact about the sphere's topology, holding whether or not anyone regards the point as sacred; a purely secular surveying problem, finding the bearing to any fixed reference point from anywhere on Earth, inherits the identical hairy-ball obstruction with no religious content at all, as Section 7.2 already made explicit. Eliade's axis mundi is not, in the first instance, a claim about any such point's geodetic coordinates; it is a claim about what makes space experientially navigable in the first place, prior to and independent of any particular center's location, and on his account a sacred center is less discovered at a pre-existing geometric point than constitutive of the very orientation within which 'pointing toward it' becomes a meaningful act for the community in question. Two

different traditions could, on Eliade's reading, each organize space around their own axis mundi without either being geodetically mistaken about the other's location, because the disagreement, if there is one, is not first a disagreement about coordinates.

Holding the two centers apart, rather than letting Eliade's vocabulary quietly annex Chapter 7's theorem or the theorem quietly flatten Eliade's phenomenology into geodesy, is this chapter's closing instance of the discipline Section 8.1 opened with: a distinction the geometry alone cannot make. The qibla is unusual, and unusually well suited to this volume's argument, precisely because it is both at once, a geodetically exact point subject to Chapter 5's topological obstruction and, for the tradition that orients toward it, something closer to Eliade's constitutive center than to an arbitrary coordinate a surveyor might equally have chosen; the companion volume's astronomers, discussed in Section 8.2, spent centuries refining the first fact precisely because the second fact made the refinement matter, and neither fact, on this volume's own terms, explains the other away.

8.8 Tuan's Topophilia: The Affective Register Orientation Leaves Out

Sections 8.1 through 8.8 have examined orientation toward a sacred center as a problem with two faces, an astronomically exact geometric question and a jurisprudential and phenomenological one, Eliade's axis mundi supplying, in Section 8.7, an ontological register in which the sacred center is not merely located but constitutive of inhabitable space itself. Yi-Fu Tuan's Topophilia, published in 1974 and foundational to humanistic geography as a discipline, approaches the bond between a person and a place from a register adjacent to Eliade's but distinct from it: not the ontological structuring of space around a sacred center, but the affective attachment, love of place in Tuan's own coinage, that develops toward particular, often unremarkable, locations through lived familiarity, sensory memory, and accumulated experience, independently of whether that place occupies any privileged position in a cosmology.

The distinction matters for this chapter because the qibla problem, as Sections 8.1 and 8.2 posed it, is a facing-problem, a question of direction toward a fixed point answerable, in principle, by geometry or jurisprudence; topophilia is not a facing-problem at all, since the attachment Tuan describes does not obviously decompose into a direction and a distance, and a person's topophilic bond to a home village is not diminished by that person's uncertainty about its bearing from anywhere else. Where Eliade's axis mundi supplies a center that organizes space cosmologically whether or not any individual feels anything toward it, Tuan's topophilia supplies an attachment that can attach to any place, sacred or entirely secular, and that requires no cosmological center to organize it; the two are compatible, a sacred center can also be topophilically loved, and De Certeau's practiced, walked space from Section 8.8 is closer kin to Tuan's register than to Eliade's, since both concern how a place is actually inhabited rather than how it is cosmologically positioned.

What topophilia adds to this chapter's closing argument is a caution against treating orientation, in either its astronomical or its jurisprudential sense, as though it exhausted what a sacred or a home place means to the person oriented toward it. Section 8.4's account of rigor as an offering rather than a requirement already gestured at something beyond mere geometric correctness in the worshipper's effort; Tuan's vocabulary makes explicit what that something is, an affective bond irreducible to bearing and distance, present in the pilgrim's relation to Mecca as much as in the resident's relation to an unremarkable home street, and not a further fact Chapter 5's or Chapter 8's geometry could in principle capture by being made more exact.

8.9 GPS and the Retreat of Embodied Dead Reckoning

Section 8.5 discussed the compass app as a return of exactness for a single, occasional task, finding the qibla bearing at a point in time and then following it by ordinary embodied means. Satellite positioning, in its now-ordinary form as continuous turn-by-turn navigation, extends the same technological logic to wayfinding in general, and does so in a way that changes more than 8.5's case,

since it can replace not just the calculation of a single bearing but the entire wayfinding activity, step by step, that Sections 7.5, 7.7, and 7.8 described as independently adequate local solutions to Chapter 5's obstruction.

The asymmetry with Section 8.5 is worth stating precisely. A qibla compass app is consulted once, and the walking or standing that follows remains an ordinary embodied competence; continuous GPS guidance is consulted throughout, so that the competence Lynch's subjects built through repeated experience (Section 7.7) may have less occasion to develop in habitual users of it. Behavioral research comparing London taxi drivers, who train for years on The Knowledge, against GPS-assisted control groups, has reported associated differences in hippocampal engagement during navigation tasks; the finding is noted here cautiously, as suggestive and still contested in its details, not as a settled result this volume can lean on the way it leans on Chapters 4 and 5's theorems.

The stakes are kept modest, in keeping with this Part's discipline. Nothing here claims that a competence lost to GPS is metaphysically essential, since 7.5, 7.7, and 7.8 already established that multiple, independently adequate solutions to the same underlying problem have coexisted across unrelated traditions and eras. The more limited observation is that a single technology can retire a whole family of distributed, local, embodied solutions at once, concentrating what had been a widely held skill into one continuously computed field — the closest the ordinary, non-astronomer's world comes to an actual instance of the global field Chapter 5 proves cannot exist as a mathematical object, existing here only as a computed and continuously updated approximation to one, not as the impossible thing itself.

8.10 Al-Bīrūnī's Spherical Solution: Exactness Before the Astronomers' Consensus

Section 8.3 described what jurisprudence did with the astronomers' exactness, treating the existence of a rigorous mathematical answer to the qibla problem as background to the legal and religious accommodation of approximate methods.

It is worth naming the specific mathematical achievement that background rests on: Abu Rayhan al-Bīrūnī's eleventh-century treatise Tahdīd Nihāyāt al-Amākin, which solves the qibla problem as a genuine spherical trigonometry exercise, computing the great-circle direction from a given city to Mecca using the coordinates of both points and the spherical law of cosines, centuries before the general retroazimuthal projections Chapter 7 discusses were formalized as such.

Al-Bīrūnī's method is worth distinguishing carefully from the later, more general retroazimuthal constructions Section 7.3 introduces. He does not produce a projection at all, in the sense Chapter 4 gives that word — no map on which every point shows its own true bearing to Mecca simultaneously — but a pointwise calculation, solved city by city as needed, using the same underlying spherical geometry a retroazimuthal projection would later encode graphically for every point at once. The distinction matters because it shows the two singularities Section 7.3 identifies, the center and the antipode, were already implicit in a working computational method centuries before anyone asked whether a single picture could display the answer everywhere without exception.

This also sharpens Section 8.4's claim that rigor functioned as an offering rather than a requirement. Al-Bīrūnī's calculation was available, exact, and largely unused for ordinary mosque orientation, which continued to rely on the simpler cardinal-direction and local-tradition methods Section 8.2 describes; the gap between having an exact answer and requiring one, which Section 8.4 treats as a considered accommodation rather than an oversight, is not a modern phenomenon introduced by compass apps (Section 8.5) but is visible already in the readiness with which al-Bīrūnī's contemporaries let his solved problem remain a specialist's tool rather than a universal standard.

8.11 A Case in Point: Praying Toward Mecca From the International Space Station

Section 8.4 read juristic tolerance of approximate qibla-facing as rigor treated as an offering rather than a strict requirement, and Section 8.10 showed Al-Bīrūnī reaching for genuine spherical exactness centuries before consensus demanded it. A single, recent, and unusually well-documented case shows the same jurisprudential structure extended to a situation neither the medieval astronomers nor Section 8.2's two ordinary methods were built to handle. In 2006, Malaysia's national space agency convened a two-day conference of roughly one hundred fifty Islamic scholars and scientists to determine how a Muslim astronaut should pray in orbit, in anticipation of Sheikh Muszaphar Shukor's flight to the International Space Station the following October, a flight that fell during Ramadan. The resulting document, approved by Malaysia's National Fatwa Council and titled a guideline for performing Islamic worship aboard the station, confronted a version of the qibla problem genuinely harder than any this chapter has yet considered: the International Space Station orbits the Earth roughly every ninety minutes at a speed near 27,600 kilometers per hour, so that the compass bearing toward Mecca, well-defined and slowly varying for any earthbound traveler, can swing through nearly one hundred eighty degrees over the course of a single prayer, a rate of angular change no jurisprudential tradition built around travelers, ships, or even aircraft had ever needed to accommodate.

The guideline's resolution follows precisely the pattern Section 8.4 identified in the terrestrial case, tolerance calibrated to genuine difficulty rather than proof-driven exactness, but adds a construction the terrestrial tradition never required: astronauts are instructed to face the qibla according to their capability, in a stated order of preference, and one endorsed method substitutes for a literal compass bearing a fixed geometric abstraction, the direction obtained by projecting the Kaaba straight outward from the Earth's surface to the station's own orbital altitude, so that the astronaut faces not Mecca itself but a point in space vertically above it at the moment of prayer. This is not a solution Section 8.3's astronomers or Section 8.10's Al-Bīrūnī could have anticipated, since it presupposes an orbital, not terrestrial, frame of reference, but it is recognizably

the same move: where exactness becomes genuinely unrecoverable, a jurisprudential tradition already comfortable, per Section 8.4, with treating rigor as an offering rather than a requirement, supplies a principled substitute rather than either abandoning the practice or pretending the difficulty away, extending a centuries-old accommodation to a case its original framers could not have imagined needing to accommodate.

Chapter 9 — Datum, Epoch, Drift: Geodetic Stability as an Ontological Problem

9.1 The ground was never the question, until it was

Chapters 7 and 8 shared an assumption worth making explicit now that it is about to be dropped: that the sphere itself — the surface on which C, P, and every other locality sit — is fixed. Locate P, compute the bearing to C, and the answer is good until either point moves. Historically, before satellite geodesy, this assumption was not a simplification but simply invisible; the earth's crust moves at rates no pre-modern instrument could have detected against the noise of its own measurement error, and every practical purpose the retroazimuthal problem served could treat the earth's surface as a rigid sphere without loss. What the companion volume The New Datum documents, in its own technical terms, is the point at which this invisibility ended: continental crust turns out to drift, continuously and measurably, at rates of a few centimeters a year — small in a human lifetime, but large enough, accumulated over the decades a geodetic reference frame is expected to serve, to move a fixed point's coordinates by a distance no longer inside anyone's error budget. The North American plate, on which the older reference frame NAD83 was in effect anchored, is one such moving surface; NATRF2022, its replacement, is the response.

9.2 A coordinate is a promise with a date attached

What NATRF2022 changes, philosophically, is not the method of measurement but the logical form of the answer a datum gives. A coordinate stated against a rigid, non-moving reference frame is a timeless proposition: it is simply true, or false, that this point has these coordinates, independent of when the question is asked. A coordinate stated against a reference frame that is itself understood to be attached to a drifting plate is, by contrast, true only relative to a specified epoch — a stated moment at which the frame's own position is fixed by convention, against which later or earlier measurements must be corrected by a known velocity before they can be compared. This is a small change in bookkeeping with a large change in what a coordinate is: no longer a location's timeless address, but a claim indexed essentially to a date, which is false, or at least misleading, the moment that date is dropped from it. The companion volume treats this technically as a matter of maintaining a time-dependent frame and a published velocity model; what should be noticed here is that the fix does not restore the older, dateless notion of position. It replaces it with a strictly more complicated one, in which stability has to be re-derived at each use from a stated epoch and a rate of change, rather than simply assumed.

9.3 Two kinds of unsteadiness

It is worth setting this explicitly against Chapter 8's problem, because the two look superficially similar — both concern a fact that can no longer be treated as simply fixed — and are not. The qibla problem left C fixed and asked about the adequacy of P's knowledge of the bearing to it; the difficulty lived entirely on the side of the method, not the geometry, which supplied a single true answer at every moment even when the method available to compute it was only approximate. The datum problem is different in kind: it is not that our knowledge of a fixed fact is imperfect, but that the fact itself — the coordinates of the reference points the whole frame is built from — does not stay put long enough to be simply known once and then relied upon. Chapter 7 distinguished two singularities, the center's completion and the antipode's genuine loss of orientation; the datum problem introduces a third kind of difficulty, orthogonal

to both, in which the trouble is not located at any particular point on the sphere at all but in the sphere's own motion beneath every point at once. There is no antipode to the datum problem, and no arrival that resolves it, because the thing in motion is the frame itself, not anyone's position within it.

9.4 What Part III has now shown, together

Read together, Chapters 7 through 9 trace a single widening of the same question. Chapter 7 asked what it is to face a fixed elsewhere from a fixed here; Chapter 8 asked how much exactness that facing needs, given who is doing it and why; Chapter 9 asks what happens when 'here' itself stops being a fact one can simply state without a date on it. None of the three chapters weakens the mathematics that grounds them — the hairy ball obstruction, the trigonometry of bearing, and the velocity model of a drifting plate are exactly as rigorous as their companion volumes establish them to be. What changes, moving through the three, is how much of ordinary human orientation those results turn out to touch: not only the sailor's or the worshipper's momentary question of which way to turn, but the geodesist's much larger question of whether the ground being turned upon can be counted on to hold still while the question is asked. Part IV takes up what all three cases together suggest about representation in general, beginning from the most general form of the lesson: that a map's distortion, like a datum's drift, is not a failure to be corrected away but a condition that any adequate philosophy of representation has to build in from the start.

9.5 Not Only Plates: Sea Level, Ice, and the Epoch Problem Multiplied

Section 9.1 introduced the datum problem through horizontal drift, the slow lateral motion of tectonic plates that NATRF2022 accounts for with a time-dependent frame and a published velocity model. Horizontal position is not, however, the only coordinate the ground supplies, and it is worth noting, briefly, that the vertical coordinate, elevation, height above a reference surface

commonly tied to mean sea level, turns out to raise the same epoch problem in a form if anything less forgiving, because more of what moves the reference surface is itself currently accelerating. A geoid-based vertical datum is defined relative to a modeled equipotential surface approximating mean sea level, and that surface is not fixed either: thermal expansion of ocean water and the addition of meltwater from land ice raise it globally, while the ground beneath a given location can itself be rising or subsiding, through glacial isostatic adjustment, the slow rebound of land once depressed by ice-age glaciers and still relaxing upward thousands of years later, or through groundwater extraction and sediment compaction, which lower it, sometimes rapidly, in exactly the low-lying coastal regions where the distance to the reference surface matters most.

The philosophical point Section 9.2 drew from horizontal drift, that a coordinate under an honest datum is a claim indexed to an epoch, not a location's timeless address, applies here with an added complication the horizontal case does not have: land and sea are moving relative to each other from both sides at once, at rates that are not, unlike plate motion, well approximated by a single steady velocity extrapolated indefinitely into the future, because both the thermal and the ice-melt contributions to sea-level rise are themselves changing over the modeling horizon a modern vertical datum is meant to serve. Where NATRF2022's horizontal frame can publish a velocity and expect it to remain a reasonably good linear model for years, a vertical datum serving the same period faces a reference surface whose own rate of change is one of the more actively revised quantities in earth science. This does not weaken Section 9.2's argument; if anything it strengthens it, by showing that the 'coordinate as a promise with a date attached' is not a peculiarity of horizontal plate tectonics that a sufficiently patient geodesist might eventually fix with a good enough velocity model, but a structural feature of any datum built on a reference surface that is itself a physical object, subject to physical processes that do not hold still merely because a coordinate system needs them to.

9.6 The Datum as an Object Extended in Time: A Metaphysical Aside

Section 9.2 described what changes, philosophically, when a coordinate is understood as indexed to an epoch rather than stated timelessly: a shift, in the vocabulary of the metaphysics of persistence, from treating the reference frame as a single unchanging thing wholly present at every moment it is used, to treating it as a thing whose different temporal stages must be related by an explicit rule — NATRF2022's velocity model — before a claim made at one stage can be compared with a claim made at another. This is, without needing the technical apparatus of the debate to be imported wholesale, recognizably the distinction between endurantism, the view that an object persists through time by being wholly present at each moment of its existence, numerically the same thing at every stage, and perdurantism, or four-dimensionalism, the view that a persisting object is better understood as extended through time the way it is extended through space, so that what exists at any one moment is only a temporal part, a slice, of a larger four-dimensional whole, and identity over time is a matter of those slices being appropriately related rather than a single thing's remaining wholly present throughout.

The reference frame NATRF2022 replaces was used, in practice, as an endurantist object: 'the coordinates of this benchmark' functioned as a fact wholly available at any moment of use, unindexed to when the question was asked, because the plate's motion was too slow to register against any practical error budget. What NATRF2022 makes explicit is that the frame was never actually an endurantist object in the metaphysically stricter sense; it only behaved like one at the resolution earlier instruments could achieve. Once resolution improves enough to detect the drift, the honest description becomes perdurantist almost by necessity: a coordinate is now a claim about one temporal slice of a four-dimensional frame, related to other slices by the published velocity model, and 'the' position of a benchmark, asked without a date, is no longer even well-formed as a question, the same way asking for the perdurantist's whole four-dimensional object's location at no particular time is not well-formed. This is offered, in keeping with Section 9.2's own caution about the difference between a

rigorous constraint and a useful vocabulary, only as a vocabulary — the metaphysics of persistence does not supply a further proof binding the datum problem, which remains exactly and only what the companion volume's geodesy establishes; it supplies a name, already worked out in a wholly different philosophical literature, for the shape of the change NATRF2022 makes precise.

9.7 Husserl's Ark That Does Not Move, and the Ground That Does

Chapters 7 through 9 have kept two registers apart throughout, following the Preliminary Note's discipline: the geometry of facing a fixed center (Chapter 7), the historical practice of computing that facing to a stated tolerance (Chapter 8), and the strictly technical sense in which the reference frame itself must now be dated (Chapter 9). A text from a wholly different tradition states, with unexpected precision, the phenomenological mirror image of exactly this sequence, and closing Part III with it earns its place only if the mirroring is stated exactly rather than gestured at. Edmund Husserl's manuscript usually translated as 'The Earth Does Not Move,' written in 1934 and unpublished in his lifetime, argues that the Earth functions, in lived spatial experience, not as one moving body among others but as the 'ark,' the constitutive ground zero relative to which motion and rest are given their sense in the first place; for the Earth itself to be experienced as moving would require a further, more encompassing ground against which to measure that motion, and no such further ground is available to embodied perceptual life. On this analysis, the Copernican fact that the Earth orbits the sun is not phenomenologically false, but it belongs to a different, theoretically mediated register from the one in which 'here,' the standing place from which every other motion is judged, is originarily given as at rest.

Chapter 9's datum problem states, in the register of technical geodesy, almost the reverse claim, and the reversal is instructive rather than embarrassing to either side. NATRF2022 exists precisely because the ground, treated as Husserl's unmoving ark by nearly every user of a coordinate for nearly the whole history of surveying, does move, measurably, continuously, and by an amount that now

exceeds the error budget the ark-status was quietly relying on; Section 9.2's 'coordinate as a promise with a date attached' is the mathematical bookkeeping required exactly because the earth, at geodetic resolution, has stopped behaving like Husserl's ark and started behaving like one moving body among others, tracked against a reference frame external to it. This is not a refutation of Husserl's phenomenology by geodesy, nor of NATRF2022 by phenomenology; the two claims are about different senses of 'the ground moves' — the constitutive sense in which the Earth is the un-thematized zero-point that makes any experience of motion possible at all, and the empirical sense in which a specific plate's specific velocity, in millimeters per year, is measured against an external celestial reference frame maintained by very long baseline interferometry. Husserl's ark can remain, correctly, the phenomenological zero-point of ordinary embodied orientation, the ground Chapter 7's pilgrim stands on without needing to date it, while the geodesist's plate, measured in a register Husserl's analysis was never about, drifts beneath that same ground at a rate now large enough to matter for anyone building a reference frame meant to outlast a single human lifetime.

What Part III has shown, across Chapters 7 through 9 and now this closing pair with Husserl, is a single structure recurring at three depths: a fixed point that orientation is directed toward (Chapter 7), a tolerance within which that direction must be known (Chapter 8), and a ground whose own fixity turns out to be a matter of resolution and register rather than an absolute fact (Chapter 9, sharpened here). Husserl's ark is not a fourth technical case alongside the other three; it is the reminder, owed to the reader at the close of a Part built entirely on technical results, of why any of this was ever worth wanting to know in the first place — because the ground one stands on, however it is finally described, is the one thing orientation cannot do without.

9.8 Bergson's Duration and the Datum's Photographic Instant

Section 9.6 treated the datum as an object extended in time, and Section 9.7 read Husserl on the ark that does not move against the ground that does. A further

phenomenological resource, closer to the specific texture of Chapter 9's problem than either, is Henri Bergson's distinction between duration (durée), lived time as a continuous, indivisible flow, and the spatialized, clock-segmented time that science and ordinary measurement substitute for it, treating a flowing process as a series of discrete instants strung along a line.

A geodetic datum, on the account Section 9.2 already gave of a coordinate as a promise with a date attached, is exactly this kind of spatialization applied to the ground itself: NATRF2022, or any epoch-referenced datum, freezes a continuously drifting surface at one named instant, treating a process (plate motion, ongoing at every moment) as if it were a fixed background state photographed once. Bergson's critique of spatialized time, developed originally against a wholly different target — Zeno's paradoxes and the physicist's treatment of motion as a sequence of positions — describes with some precision what an epoch does to geological time: not a falsification, since Section 9.3 already distinguished a useful idealization from an error, but a necessary substitution of the photographable for the flowing.

The limit of the borrowing matters here as much as the borrowing itself. Bergson held that spatialized time falsifies duration whenever it is mistaken for the thing itself, a strong and contested thesis about consciousness that this volume takes no position on; what transfers cleanly to geodesy is only the narrower structural point, that any epoch-referenced datum is, by its very construction, a spatialization of a process that does not itself pause at the epoch date — a description of what NATRF2022 does, not an argument, borrowed from Bergson or anyone else, about whether it does something illegitimate in doing it.

9.9 Wegener's Reception: The Sociology of Accepting a Moving Ground

Sections 9.1 through 9.3 established that the ground was never geodetically stationary, only assumed to be for practical purposes until continental drift and later plate tectonics made the assumption untenable. The history of how that finding was received is worth adding to Chapter 9's account, because it

complicates any simple story in which better evidence straightforwardly produces belief.

Alfred Wegener proposed continental drift in 1912, backed by a battery of evidence Section 9.1 could equally well have opened with — the jigsaw fit of the South American and African coastlines, matching fossil distributions and rock strata across the Atlantic, matching glacial striations. The proposal was rejected by most geologists for roughly four decades, not primarily for lack of evidence but for lack of a plausible mechanism: Wegener, an outsider to geology trained in meteorology, could not say what force moved continents through solid oceanic rock, and this absence, rather than any specific counter-evidence, anchored the resistance until seafloor spreading supplied the missing mechanism in the 1960s.

The episode bears usefully on Section 6.5's distinction between underdetermination and relativism, applied here to acceptance rather than to choice among projections. The geological community's decades-long resistance was not, on the balance of the historical record, an irrational clinging to convention against decisive proof; it reflected a defensible methodological demand — evidence for a pattern is not yet evidence for a mechanism — that happened, in this instance, to delay recognition of something Chapter 9 treats as settled fact. The datum's drift, whose reality Section 9.3 takes for granted, was itself received exactly as slowly and as reasonably contested as any other major revision to a background assumption long treated as safely fixed.

9.10 When the Ground Moves Under a Property Line: A Cadastral Postscript

Sections 9.1 through 9.3 treated the datum problem at the level of the reference frame itself: what a coordinate is, once the surface it is stated against is understood to drift. It is worth closing the chapter by noting, briefly, an institutional consequence Chapters 7 and 8 did not have an analogue for, because the qibla problem concerns a direction and the datum problem concerns a fixed record of location that other records, and other people's legal entitlements, are

built on top of. A cadastral boundary, the legally recorded line separating one parcel of land from its neighbor, is ordinarily surveyed once and recorded as a fixed fact, coordinates and bearings fixed at the moment of survey and expected to remain authoritative until a new survey is ordered for some unrelated reason. Section 9.2's argument that a coordinate under an honest datum is a promise indexed to an epoch, not a location's timeless address, applies to this record exactly as it applies to any other, and with a consequence the geodesist's own use of the frame does not carry: two adjacent parcels, surveyed years apart against reference frames whose underlying epochs differ, can each be correctly stated relative to its own epoch and still disagree, on the ground, about where the shared boundary between them actually falls.

This is not a hypothetical difficulty invented for the chapter's sake; national geodetic agencies undertaking a datum modernization of the kind the companion volume documents routinely publish guidance for surveyors and land administrators on exactly this transition, because a boundary that was unambiguous under the older, effectively dateless convention becomes, under the newer and more honest one, a claim that must specify which epoch's realization of the frame it is stated in before two independently correct-seeming surveys can even be compared. The philosophical point is a specific instance of Section 9.4's wider claim, that what changes across Chapters 7 through 9 is how much of ordinary practice the underlying mathematics turns out to touch: the hairy ball theorem's two singularities are a fact about bearing that a pilgrim or a sailor must accommodate; the datum's drift is a fact about location that a title office, a fence line, and a tax assessment must accommodate as well, whether or not anyone administering them has ever heard of NATRF2022 or the theorem behind it.

The chapter's discipline, stated in Section 9.2 and held to since, applies here without modification: this section reports a real administrative consequence of a real technical transition, sourced to the same companion volume's treatment of datum modernization, not a speculative extrapolation of what might happen to property law somewhere. What it adds to Chapter 9's argument is only the

observation that a philosophy of representation content to stop at the geodesist's reference frame would miss exactly the point Chapter 13 will later insist on for maps in general — that a representational convention, however well-founded technically, is never merely descriptive once other institutions have built binding commitments on top of it.

9.11 A Case in Point: The Pacific Islands Forum and the Refusal to Let Baselines Ambulate

Section 9.2 described a coordinate as a promise with a date attached, and Section 9.10 traced what happens when a cadastral boundary, fixed by survey at one epoch, is asked to remain meaningful as the ground beneath it drifts. A more consequential instance of the same structure, adjudicated not by a local land registry but by international law, makes the point at the scale of entire nations rather than individual parcels. Under the ordinary, so-called ambulatory doctrine of the law of the sea, a coastal state's maritime baseline, and the territorial sea, exclusive economic zone, and continental shelf claims measured outward from it, is understood to move as the physical coastline itself moves, so that a receding shoreline caused by erosion or, increasingly, by climate-driven sea-level rise, would in principle shrink a state's maritime entitlements along with its land. On 6 August 2021, at the Fifty-First Pacific Islands Forum, the leaders of eighteen Pacific states declared their intention to do exactly what Section 9.2's language anticipates: to treat their maritime baselines, once established and deposited with the United Nations Secretary-General under the 1982 Law of the Sea Convention, as fixed at that epoch, and to decline to update them regardless of subsequent, climate-driven physical change to the coastlines from which they were originally measured.

The declaration is worth reading against this chapter's vocabulary rather than only against its own legal one, because it is, in the terms Section 9.2 introduced, an explicit refusal to let a geodetic promise expire when the ground it was made about changes: where NATRF2022, Appendix E's case study, responds to instability by re-issuing the promise at a new epoch, the Pacific Islands Forum's

declaration responds to a structurally identical instability, a physical baseline drifting out from under a legal coordinate, by insisting the old promise be honored regardless of the drift. Nothing in this volume's proofs adjudicates whether that legal choice is wise, defensible under the Convention's own terms, or likely to be accepted by states with competing interests in the same waters, questions properly belonging to international law and well outside this volume's competence; what belongs here is only the observation that the epoch problem Section 9.5 generalized beyond plate tectonics to sea level and ice mass has now generated its own concrete instance of the choice Section 9.2 described in the abstract, between letting a coordinate ambulate with a changing ground and freezing it by declaration, at a scale, and with stakes, far beyond the cadastral postscript with which Section 9.10 first illustrated the same dilemma.

PART IV

Distortion as a Condition of Representation

Chapter 10 — ‘Every Map Lies’: Distortion as Condition, Not Defect

10.1 Collecting the theorems

Recall three impossibility results established with full rigor in the companion volumes and used in this one without re-derivation: Tissot's theorem (Chapter 4) that no projection preserves both angle and area everywhere; the hairy ball theorem (Chapter 5) that no continuous, nonvanishing tangent vector field exists on the sphere, hence no everywhere-consistent orientation field; and the datum problem (Chapter 9) that even the reference frame itself cannot be held fixed once time is taken seriously. Despite covering different mathematical territory — the differential geometry of surfaces, algebraic topology, geodynamics — the three share a structure worth stating plainly: each is a genuine 'cannot', not a 'not yet'. No amount of ingenuity within the terms of the problem removes it; ingenuity only changes where the unavoidable cost is paid, or how the residual freedom is used, a point already the substance of Chapter 6's conventionalism, Chapter 7's asymmetric singularities, and Chapter 8's graded jurisprudential standard.

10.2 From cartographic distortion to representational distortion in general

The cartographic case is unusually clean because its 'cannot' is a theorem, not an empirical generalization. Is there anything cartography-specific keeping the lesson from applying wherever a lower-dimensional or otherwise reduced representation stands in for a richer structure? Chapter 3 already argued no: what makes cartography a good test case is precisely that its impossibility is provable, not that its structure is unique. A scientific model that must be simpler than the system it models, a legal category that must sort continuously varying cases into discrete classes, a verbal description that must select some features of an experience and drop others — all share with cartography the basic shape of

reduction under constraint. This is not the claim that these other domains have their own Tissot-style theorems; most do not, and where a domain-specific impossibility proof does exist, it deserves the same rigor demanded of the cartographic case, not a borrowed one. It is the claim that cartography demonstrates, in a domain where the claim can be proven, a structural feature that the same reduction-under-constraint situation makes plausible elsewhere even without proof: distortion is not contingent on getting the technique wrong.

10.3 The difference between a flaw and a cost

It matters, philosophically, whether one names the residue of a reduction a 'flaw' or a 'cost'. A flaw invites a search for the better version that lacks it — the perfect projection eventually to be found, the more sophisticated model eventually to eliminate the simplification, the more careful description eventually to capture everything. A cost invites a different question: given that some cost is unavoidable, which one, paid where, is worth paying for which benefit. Chapter 4 already made this choice explicit for projections — equal-area versus conformal versus a compromise, chosen for purpose, not toward a nonexistent ideal — and the present chapter's claim is that this is the correct general attitude toward any representation under a genuine impossibility result, not a concession specific to mapmaking. Treating a mathematically unavoidable distortion as a flaw to be engineered away is not merely optimistic; once the impossibility is proven, it is a category error, mistaking a theorem for a technical shortfall.

10.4 What this does not license

The same discipline that requires taking the theorems seriously requires refusing to let them travel further than they can support. It does not follow, from the fact that all maps distort, that all distortions are equally distorting, equally motivated, or equally excusable — a point Chapter 13 takes up directly, because it is exactly where the philosophical use of these results is most often abused. Nor does it follow that because representation is always partial, no

representation is ever more accurate than another, or that accuracy itself is somehow illusory; Tissot's theorem specifies precisely how and where a given projection lies, which is the opposite of denying that some projections lie less, or more usefully, than others for a given task. The impossibility results license humility about completeness, not skepticism about better and worse. Holding both without collapsing one into the other is the discipline this volume has tried to practice from its opening page, and it is the discipline Chapter 13 needs most, because letting a mathematical impossibility do political work it was never built for is exactly where the two temptations named in Chapter 1 return in their most consequential form.

10.5 A Test Case from Outside Cartography: The Bias-Variance Trade-off

Section 10.2 argued, carefully marking the claim as analogy rather than theorem, that other representational domains plausibly share cartography's basic shape of reduction under constraint, without claiming any of them has its own Tissot-style proof. It is worth working one candidate case through in enough detail to show what the analogy does and does not license, because the discipline Section 10.4 insists on is easiest to state in the abstract and easiest to lose in practice. Statistical and machine-learning models that fit data face a well-known trade-off between bias, the systematic error a model makes by being too simple to capture the true pattern, and variance, the error it makes by being so flexible that it fits the noise in a particular sample rather than the pattern the sample was drawn to reveal; a model simple enough to generalize well typically underfits, a model flexible enough to fit the training data closely typically overfits, and a large body of statistical theory formalizes exactly how these two error sources trade against each other as a model's complexity is tuned. The surface resemblance to Chapter 4's Tissot trade-off, angle against area there, bias against variance here, with a family of intermediate compromises available in both cases, is close enough that the analogy suggests itself immediately, which is precisely why it needs testing rather than simply enjoying.

Tested, the analogy holds up better than many candidates would, but not perfectly, and the difference is instructive. Like Tissot's theorem, the bias-variance decomposition is a proven mathematical result, not a loose empirical generalization, it follows from the algebra of expected squared error, and in that narrow sense the statistical case is not merely image or analogy but a second, independent instance of a proven representational impossibility, closer to Part II's own register than to the merely suggestive comparisons Section 10.2 mostly has in mind. Where the analogy runs out is in what the trade-off is a trade-off between: Tissot's theorem fixes, once and for all, that no projection anywhere achieves both properties, for any surface with the sphere's curvature; the bias-variance decomposition instead describes a relationship between a specific model class and a specific data-generating process, one that changes as more data becomes available or as a better-suited model class is chosen, in a way Tissot's theorem, tied to the fixed geometric fact of curvature, does not. The statistical trade-off is real and provable, but it is a fact about the fit between a family of models and a problem, not, like Tissot's, a fact about representation as such that no choice of family could ever escape, a distinction worth holding onto exactly because the two trade-offs otherwise look so alike.

10.6 Idealization Taxonomies: Weisberg's Three Kinds, and Cartwright's Nomological Machines

Section 10.5 compared the projection trade-off to the bias–variance trade-off of statistical learning and marked carefully where the analogy held and where it did not. A further and more direct comparison is available in the philosophy of science's own vocabulary for idealization, and it is worth making before Chapters 11 and 12 take up the wider question of scientific models directly. Weisberg distinguishes three kinds of idealization. Galilean idealization simplifies a model for present tractability while intending, at least in principle, to reintroduce the stripped-away complexity later, so that the idealized model is a way-station toward a fuller one rather than a terminus. Minimalist idealization strips away every causal factor except the one or few that actually make the

relevant difference to an outcome, and does so with no intention of ever adding the rest back, because the rest genuinely does not matter to the pattern being explained. Multiple-models idealization abandons the search for one unified model altogether and instead builds several incompatible simplified models of the same target, each capturing a different facet, deliberately declining to reconcile them into a single picture.

Cartography's relationship to this taxonomy is precise rather than approximate, which is unusual enough to be worth stating plainly. It is not Galilean idealization: Galilean idealization promises, at least in principle, a path back to the full target, and Chapter 4's theorem forecloses that path outright — there is no more complete projection waiting to be reintroduced, only the proof that none can exist. It is not minimalist idealization either, in Weisberg's causal sense, since no single projection isolates 'the' causally relevant factor governing shape while discarding causally irrelevant ones; angle and area are not in a causal relation of that kind, but a joint geometric constraint. What cartography's actual practice instantiates, exactly rather than loosely, is multiple-models idealization: Section 4.4's case studies — Mercator for bearing, Albers for area, the compromise projections for neither exactly — are precisely a suite of incompatible models of the same target, each preserving a different structure, used side by side without any attempt at reconciliation into one map that does everything. This is one of the closest transfers this volume can honestly claim, and the reason it is close is worth stating rather than left implicit: elsewhere, multiple-models idealization is typically adopted because the best available models happen, as a matter of current scientific practice, to be incompatible; in cartography the same arrangement follows from a proof that no reconciled alternative could exist even in principle. The structure is identical; its modal status is not.

Cartwright's account of nomological machines and the accompanying 'simulacrum' view of models, from How the Laws of Physics Lie, offers a different and more radical comparison. A nomological machine, on Cartwright's account, is a fixed arrangement of components, real or idealized, shielded from interference, built specifically to let a particular law operate in a clean and

repeatable way — closer to an engineered laboratory setup than to a picture of the world as it already is. Cartwright presses this into a strong thesis: that most physical laws, read literally, are true only of the models built to house them and not of the world outside those models, which she calls a simulacrum rather than a representation in any realist sense. This resonates with Section 4.3's discussion of choosing among the surviving, mutually incompatible projections after Tissot's theorem forecloses the ideal one, and with Section 6.2's conventionalism about which of several equally consistent geometries to adopt. But the resonance should not be allowed to blur an important asymmetry this volume's own position does not share. Nothing in Tissot's theorem or the hairy ball theorem requires denying that a projection can be honestly evaluated as more or less accurate in Section 3.2's third, purpose-relative sense, and nothing in them requires denying that the geographic facts a map depicts are objectively true independent of any map at all. Cartwright's simulacrum account is a stronger and considerably more contested claim, amounting to a fictionalism about physical law as such; this volume borrows the structural insight about engineered, purpose-built idealization without adopting the wider anti-realism, and says so explicitly rather than letting the borrowing imply an endorsement it has not argued for.

10.7 A Third Test Case: The No Free Lunch Theorem

Section 10.5 tested one candidate analogy, the bias-variance trade-off, against the standard Section 10.4 sets, and found it a genuine second proof with a narrower scope than Tissot's own. A second candidate is worth the same treatment, because it is, if anything, an even closer surface match to the cartographic case and therefore an even more useful test of the discipline this chapter asks the reader to practice. The no free lunch theorems for optimization, proved by Wolpert and Macready, establish that, averaged uniformly over all possible optimization problems, every search algorithm performs identically — no algorithm can outperform random search when its performance is averaged across the entire space of possible objective functions, because whatever

structure lets an algorithm exploit some problems well is exactly what causes it to perform correspondingly worse on problems lacking that structure. Stated this way, the theorem has an unmistakably Tissot-shaped structure: a virtue (performance on some class of problems) trades off, provably and exactly, against its absence on the complementary class, and no algorithm escapes the trade-off by cleverness, because the theorem is proved for the full space of possible problems, not observed as a tendency across the problems actually encountered in practice.

Tested against Section 10.4's standard, the resemblance again holds up further than most candidates would, and again not all the way. Like Tissot's theorem and unlike the bias-variance case, the no free lunch result is a statement about the full space of possible targets, not about a fit between one model class and one data-generating process — in this respect it is structurally closer to Tissot's theorem, which likewise ranges over every point of every possible projection, than the bias-variance trade-off examined in Section 10.5 was. Where it diverges from Tissot's case is in what licenses the uniform average the theorem is proved over: Tissot's theorem holds for the sphere's actual, given curvature, a fact about the one surface cartography actually needs to map, while the no free lunch theorem's force depends on averaging across a hypothetical space of all possible objective functions, most of which — pure noise, uncorrelated with anything an actual scientist or engineer would ever pose as a problem worth solving — are not the kind of problem practitioners in fact encounter, which is precisely why the theorem, though true without qualification, is widely and correctly regarded in the optimization literature as a much weaker practical constraint than its initial statement suggests: real problems are not drawn uniformly from the space the theorem ranges over, and an algorithm exploiting the structure real problems actually have is not thereby refuted by a theorem about problems that do not occur. The lesson for this chapter's larger argument is exactly the disanalogy Section 10.5 already modeled: a second proven trade-off, real and exact within its stated domain, whose domain of application turns out, on the closest inspection, not to coincide with the domain a first glance at its vocabulary would suggest.

10.8 The Akaike Information Criterion: The Bias-Variance Trade-off Made Computable

Section 10.5 introduced the bias-variance trade-off as a test case for this chapter's central claim, that distortion is often a condition of representation rather than a removable defect, drawing the parallel between a projection's forced sacrifice along one of Tissot's axes and a statistical model's forced sacrifice between fitting its training data too loosely and fitting it, and its noise, too closely. Hirotugu Akaike's information criterion, published in 1974, supplies the parallel a computable instrument rather than a qualitative analogy: AIC estimates, for a candidate model, a quantity proportional to twice the number of fitted parameters minus twice the model's maximized log-likelihood, a penalty for complexity set directly against a reward for fit, derived from the expected Kullback-Leibler divergence between the candidate model and the unknown true data-generating process, and the model minimizing this quantity among the candidates considered is preferred, not because it is closest to certainly true, but because it minimizes expected information loss.

The cartographic reading is closer than a family resemblance between two trade-offs. A projection is, in this comparison, itself a kind of model of the sphere's geometry, fitted with a particular functional form, and Tissot's theorem is the demonstration that no functional form, however many free parameters it is given, drives the misfit to zero; adding parameters to a projection's construction, as the Dymaxion projection's polyhedral folding did in Section 4.5, can relocate the misfit without eliminating the total quantity Chapter 4's theorem bounds below, exactly as AIC's complexity penalty formalizes the suspicion, already present qualitatively in Section 10.5, that adding parameters to a statistical model improves apparent fit for reasons that do not track improved fit to the true generating process. What AIC contributes that the qualitative parallel could not is a stopping rule: a principled point, computable rather than merely felt, at which further complexity ceases to be justified by the corresponding gain in fit, the numerical analogue of Section 4.3's purpose-relative choice among Tissot's survivors.

The disanalogy is equally worth stating precisely, in keeping with Section 10.4's insistence that distortion-as-condition not be allowed to license indifference to particular distortions. AIC ranks models against each other relative to an unknown true distribution that is assumed, however idealized, to exist; Tissot's theorem requires no such assumption, since the sphere's Gaussian curvature is exactly known and the misfit it forces is a proven lower bound, not an estimate subject to the sampling uncertainty every AIC comparison inherits from finite data. The statistical case is, in this sense, the harder-won of the two: cartography enjoys a certainty about the source of its trade-off that model selection, working from data alone toward a generating process it can never directly inspect, does not have, which is one more reason Chapter 4's theorem has served this volume as a comparatively clean anchor for a trade-off that, everywhere else it appears, must be approximated rather than proven.

10.9 The Manifold Hypothesis: A Contemporary Instance of the Bias-Variance Family

Sections 10.5 through 10.8 assembled test cases from outside cartography — the bias-variance trade-off, idealization taxonomies, the no free lunch theorem, the Akaike information criterion — to show the cartographic trade-off's shape recurring across statistics. A further and more recent instance, closer to this volume's own opening subject matter than any of the four, comes from machine learning's manifold hypothesis: the working assumption that high-dimensional data such as natural images or text embeddings lies near a much lower-dimensional, typically curved manifold embedded within the full ambient space.

The connection to Chapter 3 is closer than a further analogy. Dimensionality-reduction techniques — t-SNE, UMAP, autoencoders — are, formally, exactly the situation Section 3.2 examined: a structure-preserving correspondence sought between a richer object, the data manifold, and a poorer one, a two- or three-dimensional visualization or embedding, for purposes of inspection or downstream use. They inherit a Tissot-shaped trade-off for the

same reason a curved surface cannot be flattened without cost: no reduction can simultaneously preserve all pairwise distances, local neighborhood structure, and global topology of a genuinely curved manifold, and practitioners speak, in language that echoes Section 4.2 without having borrowed it, of trading off local structure preservation against global structure preservation when choosing between these methods.

What earns this its own subsection rather than folding it into Section 10.5 is a difference in epistemic status worth keeping visible. The earlier statistical cases either rest on a proved theorem (Section 10.7's no free lunch result) or a mathematically explicit model family (Section 10.8's AIC); the manifold hypothesis's underlying surface is not given by physics or proved to exist but inferred from data and typically unknown in closed form, so its trade-off is diagnosed empirically — through reconstruction error, or trustworthiness and continuity metrics — rather than derived as Tissot's theorem is derived. It belongs, in Section 10.6's taxonomy, closer to an idealization defended by its results than to a theorem defended by its proof, and the difference is worth marking precisely because this volume has otherwise been careful to say, at every step, which register a given claim belongs to.

10.10 Box's Aphorism: 'All Models Are Wrong' as the Chapter's Condensed Restatement

Sections 10.1 through 10.8 assembled, one register at a time, the theorems and test cases supporting this chapter's title claim that distortion is a condition of representation rather than a defect in it. George Box's frequently quoted remark, that all models are wrong but some are useful, is worth adding not as a further test case alongside the bias-variance trade-off or the No Free Lunch theorem, but as the shortest available statement of what this chapter has spent eight sections proving in detail.

The aphorism is usually deployed loosely, as a piece of practical wisdom rather than a conclusion earned by argument, and Section 10.4 already warned against

exactly this kind of unearned generalization. What this chapter has done, and what most invocations of Box's remark have not, is show for a specific case — cartographic projection — exactly why every model must be wrong, in Chapter 4's precise sense of a provable, quantified trade-off, and exactly what 'useful' can mean once wrongness is no longer optional: not the least-wrong model, since Section 10.5's bias-variance framing shows there is no such thing independent of a stated purpose, but the model whose particular wrongness is best matched to the purpose at hand.

Read this way, Box's aphorism functions for this chapter the way a chapter epigraph might, if this volume used them: a compressed statement whose casual, quotable form conceals exactly the ambiguity Section 10.3 spent a full section disentangling, between a flaw and a cost. The aphorism says a model is wrong; it does not say, and by itself cannot say, whether that wrongness is a design choice made honestly, in Cartwright's sense (Section 12.4), or a failure to notice one. This chapter's theorems are what turn Box's slogan from something one might merely find agreeable into something one can actually use to tell the two apart.

10.11 A Case in Point: Double Descent and the Limits of the Classical Bias-Variance Story

Section 10.5 introduced the bias-variance trade-off as a test case from outside cartography, structurally parallel to Tissot's theorem without being a corollary of it: as a statistical model's complexity grows, bias falls and variance rises, and the two together trace the U-shaped risk curve every introductory course in statistical learning still teaches. That U-shape held, as an empirical regularity rather than a proven theorem, for essentially the whole history of statistical model-fitting, until a 2019 paper by Mikhail Belkin and coauthors at Ohio State and Columbia documented a second regime the classical story had never anticipated. Beyond the point at which a sufficiently large model, a deep neural network, a random forest, even certain non-neural methods, first achieves zero error on its training data, the interpolation threshold, test error does not simply continue climbing as the classical curve predicts; instead it falls a second time,

often to a level better than any model on the underparameterized side of the U ever reached, a shape Belkin's paper named double descent and showed empirically across several model families rather than only in the deep networks where the phenomenon was first noticed anecdotally.

This is worth setting beside Section 10.5's use of the bias-variance trade-off as a philosophically instructive parallel to Chapter 4's proven theorem, because it marks exactly the difference Section 10.4 warned against eliding: Tissot's theorem is a proof, true of every admissible projection without exception, while the classical bias-variance curve was always, however well-confirmed, an empirical regularity that a sufficiently different modeling regime could complicate. Double descent does not repeal the underlying trade-off, later theoretical work has largely reconciled it with a more general notion of effective, rather than nominal, model complexity, but it does show that a formal parallel this volume has been careful to call structural rather than logical can itself turn out to hold only within a bounded historical range of practice, a caution Chapter 11 will take up in its own terms when it asks how much of Tissot's proof-backed authority the looser analogies built on it are entitled to borrow.

Chapter 11 — Scientific Models and the Limits of the Analogy

11.1 Why the analogy needs its own chapter, not a paragraph

Section 10.2 raised, and deliberately declined to settle in passing, the question left open since the Preliminary Note: do scientific models distort in a sense structurally analogous to a map, or only in a sense that borrows the word? A question flagged that carefully deserves more than the paragraph that flagged it, and this chapter and the next supply what Section 3.3 promised: not an assertion of the analogy but an actual, checkable search through the literature where the analogy would have to be made good if it is to be made good at all. The test applied throughout is the one this volume has used since Chapter 1: does a given

philosophy-of-science account identify a provable, purpose-independent trade-off structurally like Tissot's theorem, or does it describe a real but practice-relative predicament that merely resembles cartography's by sharing vocabulary — model, representation, isomorphism, idealization — without sharing cartography's proof? Stated in advance, in keeping with this volume's practice of announcing its conclusions before arguing for them: no account surveyed in this chapter contains or entails a second Tissot-style theorem. That is not a disappointing result. A literature converging, independently and for its own reasons, on a predicament shaped like cartography's without ever finding cartography's proof is itself a finding, and Section 11.5 will say precisely what kind.

11.2 Partial structures and partial isomorphism: da Costa and French, Bueno

Da Costa and French's partial structures program, extended by Bueno into a general framework for scientific representation, models a theory's relation to the world using partial relations: relations that, unlike the total relations of classical model theory, need not be defined on every tuple of their domain, leaving some cases explicitly undetermined pending further inquiry. Representation, on this view, proceeds via partial isomorphism, a correspondence required to hold only on the subdomains where the relevant relations are currently defined, with the undetermined remainder neither asserted nor denied. The appeal of the framework is that it lets a working scientific theory be represented as it actually is, incomplete and under active revision, without forcing either a premature claim of full isomorphism or a skeptical denial that any representation relation holds at all.

The comparison with cartography exposes a difference in kind, not merely degree, that is worth stating without softening. A partial relation's gap is epistemic: the value is undetermined because it is not yet known, and the partial-structures framework is built to represent exactly that state of inquiry. A projection's distortion is not epistemic in this sense at all. The angular and areal

distortion at every point of a Mercator projection is not unknown, awaiting further research; it is exactly and completely computable from the projection's defining formula, a determinate fact available in closed form, and Tissot's theorem states precisely how the two kinds of distortion trade off everywhere at once. What forces cartography's incompleteness is not a gap in what is known but a proof about what can jointly be true; what forces a partial structure's incompleteness is the opposite, a gap in what is yet known about something that could, in principle, come to be fully determined by further inquiry. The word 'partial' is doing different work in the two cases, and the disanalogy is worth naming precisely because the shared word makes it easy to miss.

11.3 Giere's similarity view and the retreat from isomorphism

Section 3.5 introduced Giere's replacement of isomorphism with a graded, purpose-indexed similarity relation, and it is worth returning to that proposal here with the sharper question this chapter asks. Giere's own preferred illustration is instructive: the solar system is modeled, for different explanatory purposes, by idealizations that are mutually incompatible and that none of their users ever intended as candidates for isomorphism — a two-body Keplerian model for predicting an individual planet's position, a different, chaos-sensitive many-body model for questions about the system's long-term stability. No scientist treats either model as an attempt at a full correspondence that merely falls short; each is built, from the outset, to capture only what a specific explanatory job requires, which is genuinely close to Section 3.2's third notion of fidelity as this volume defined it for maps.

What is missing from Giere's account, and from the similarity-based literature generally, is anything resembling Chapter 4's theorem: a proof, stated once and for all, of exactly how much similarity in one respect must be traded for how much similarity in another, across the full space of possible models of a given target. Cartography can say, in closed form, precisely how conformality and equal-area preservation trade off at every point of every possible projection; nothing in Giere's framework, or in the wider modeling literature this section

surveys, offers a comparably exact law governing how much predictive accuracy about planetary position must be sacrificed for how much insight into long-term stability, or vice versa, for physical models in general. The retreat from isomorphism to similarity is common to both domains; the exact quantification of what the retreat costs is not, and Chapter 12 will show one place — Batterman's minimal models — where something closer to such a law is arguably present, though for reasons that turn out, on inspection, not to be geometric at all.

11.4 van Fraassen's constructive empiricism and the empirical substructure

Van Fraassen's constructive empiricism, already introduced in Section 3.5 as part of the semantic view, asks a theory to be empirically adequate rather than true: some substructure of its models, the part corresponding to what is observable, must be isomorphic to the appearances, while the theory is permitted, and on van Fraassen's view should be expected, to remain silent about whether the rest of the model corresponds to anything real. This is a narrowing move, and it is worth setting directly beside cartography's own narrowing move from Section 3.2: a subway map claims isomorphism, in effect, only for connectivity and sequence, remaining silent — not false, silent — about true distance and bearing, which it does not claim to preserve and does not ask its user to believe it preserves.

The structural parallel here is close enough to be worth affirming rather than merely noting, and it is one of the more successful transfers this chapter finds: both cartography's practical solution and constructive empiricism's philosophical one proceed by shrinking the scope of the isomorphism claim to exactly the substructure a specific use requires, and saying nothing beyond it. But the motivations for the narrowing are not the same kind of reason, and this volume's discipline requires saying so. Van Fraassen's narrowing is chosen, an expression of epistemic modesty about what the evidence entitles a scientist to believe concerning unobservables; a cartographer's narrowing to connectivity alone, in the subway-map case, is also a chosen design decision, but the deeper narrowing

forced by Chapter 4's theorem — that no projection of any kind can claim full isometry — is not chosen at all. It is forced by proof. Constructive empiricism narrows its claims because it judges that more would be unwarranted; cartography narrows its claims, at the deepest level, because more would be false. The two arrive at structurally similar restraint from opposite directions: modesty in one case, resignation to a demonstrated impossibility in the other.

11.5 What actually transfers from Tissot's theorem, and what does not

The discipline this volume set itself requires stating plainly, now that Sections 11.2 through 11.4 are in, exactly what has and has not been shown. What has not been shown, and was genuinely searched for rather than assumed absent: no account surveyed in this chapter — not partial structures, not Giere's similarity view, not constructive empiricism, and, Chapter 12 will confirm, not the explanatory-model literature either — contains a second theorem of Tissot's kind, an exact, provable, purpose-independent law governing how much of one representational virtue must be sacrificed for how much of another, for scientific models in general. The search was a real search, conducted according to Section 11.1's stated test, and it came back empty; that result is reported here rather than quietly dropped, in keeping with the honesty the Preliminary Note asked of every chapter.

What has been shown is different and still substantial. Three independent philosophical projects — da Costa and French's partial structures, Giere's perspectival similarity, van Fraassen's constructive empiricism — developed for their own reasons, answering their own internal questions, and not in dialogue with cartography at all, converge on describing a predicament with the same three-part shape this volume has traced through cartography with proof: first, full correspondence is neither achieved nor any longer aimed at; second, representational success is cashed out instead through a narrower relation calibrated to a specific purpose, precisely Section 3.2's third notion; third, which narrower relation is adopted is not read off the target alone but requires an additional, non-geometric — here, non-mathematical — choice, echoing Chapter

6's conventionalism. That convergence, reached independently by writers not discussing maps, is evidence, though not proof, that the shape cartography exhibits rigorously is a real and recurring shape in representational practice generally, arrived at for different reasons in each domain but recognizable across all of them. This is exactly the 'specific structural analogy, not merely a shared word' that the Preliminary Note and Section 1.2 demanded of this volume wherever it left cartography's strict terrain, achieved here by the slower and more honest route of checking rather than asserting.

11.6 Structural Realism and What Survives Theory Change

Sections 11.2 through 11.4 searched three specific accounts of scientific representation for a second Tissot-style theorem and reported, in Section 11.5, that none supplies one. A further, closely related literature deserves inclusion in the same search, because it is often invoked in exactly the borrowing-prone way Section 1.4 warned against: structural realism, in the form given it by John Worrall and developed further by James Ladyman and others, holds that across episodes of radical theory change in physics — the replacement of Fresnel's ether-based optics by Maxwell's electromagnetic theory being Worrall's own central case — what survives is not the discarded theory's ontology, the specific entities it posited, but its mathematical structure, which is typically retained, reinterpreted, in the successor theory; Fresnel's equations for reflected and refracted light survive essentially unchanged inside Maxwell's theory, even though Fresnel's ether does not survive at all. Worrall's own proposal, epistemic structural realism, holds that this pattern licenses believing in a theory's structure while remaining agnostic about the nature of the entities that structure is a structure of — a view whose vocabulary of structure surviving where entities do not sits close enough to this volume's own vocabulary of homomorphism preserving some relations while dropping others (Section 3.4) that the resemblance needs the same testing Section 11.1 applied to every other candidate.

Tested, the resemblance again proves suggestive without amounting to a second theorem. Cartography's homomorphism is a fixed relation, exactly specified in advance by a projection's defining formula, between a given surface and a given plane; structural realism's claim is instead a historical generalization about which parts of past theories have, as a matter of the historical record, tended to survive theory change, defended case by case rather than derived from a proof that structure in general must survive where ontology does not. Ladyman's own development pushes the view further, toward ontic structural realism, on which structure is not merely what we can know but, in a much stronger and more contested claim, all there ultimately is — a thesis with no cartographic analogue at all, since nothing in Tissot's theorem suggests that the sphere is nothing but the structure a projection preserves of it; the sphere's curvature, the very thing forcing the theorem, is a fact about the sphere's intrinsic geometry, prior to and independent of any projection's preserved structure, not exhausted by it. What survives the comparison, in the modest sense Section 11.5 already made this volume's standard, is a shared vocabulary of partial preservation under transformation, arrived at independently in the philosophy of science for reasons having nothing to do with maps; what does not survive is any suggestion that structural realism's historical generalization carries Tissot's theorem's exactness, provable in advance rather than defended case by case after the fact.

11.7 Wimsatt’s Robustness Analysis: Triangulating Truth Across Several False Models

Sections 11.2 through 11.6 have asked, in various technical vocabularies, how much of a scientific model’s success a single Tissot-style theorem could explain, and have generally found the transfer partial. William Wimsatt’s robustness analysis approaches the same territory from a different angle worth adding before this chapter closes, because it is built specifically for situations, like Chapter 4’s array of named projections, in which several available representations of the same target are each individually known to be false in some specified respect, and asks what, if anything, can be inferred from where

they agree. Wimsatt's proposal, developed across a series of papers on robustness in evolutionary biology and philosophy of science, is that a result invariant across multiple independent, differently-idealized models or measurement procedures earns a distinct kind of evidential credit neither model could earn alone, precisely because the different models' errors are unlikely to conspire in the same direction for unrelated reasons; convergence across independently flawed representations is evidence, though never proof, that the converging feature reflects the target rather than any one representation's particular distortion.

Chapter 4's own case studies supply an unusually clean testing ground for this idea, because the theorem governing them is known exactly rather than merely suspected, which lets Wimsatt's heuristic be checked against a ground truth the biological cases he originally worked with rarely offer. Mercator, Albers, Robinson, and the Dymaxion projection distort area, angle, and shape in different ways and by different amounts at any given location, by Chapter 4's own proof; a feature on which all four nonetheless agree, ordinal adjacency, that Canada borders the United States, that the Sahara lies between the Mediterranean coast and sub-Saharan Africa, is robust in Wimsatt's sense and, not coincidentally, is exactly the topological information Section 3.2 identified as preserved by correspondence alone, prior to any question of metric distortion. A feature on which the four projections sharply disagree, Greenland's size relative to Africa, famously inflated by Mercator and corrected by Albers or Robinson, is exactly not robust, and the disagreement itself, rather than any single projection's verdict, is the more reliable signal: it flags relative area as a quantity projection-dependent by Tissot's own theorem, not a fact any one map can be trusted to report and the others merely get wrong.

What this adds to Section 11.5's inventory of what does and does not transfer from Tissot's theorem is a usable heuristic for cases, unlike cartography, where no closed-form theorem is available to say in advance which features will be robust and which projection-dependent. Robustness analysis cannot supply Chapter 4's certainty; Wimsatt is explicit that convergence across models is

defeasible evidence, not proof, and multiple models can share a common blind spot no amount of cross-checking among them would reveal. What it supplies instead, and what makes it worth this chapter's closing attention, is a disciplined way of extending Chapter 4's lesson, that a single representation's distortions are not random but structured and knowable, into domains where no companion volume will ever prove a Tissot-style theorem outright, using multiplicity and disagreement itself as the instrument the geometric case did not need.

11.8 Mary Hesse's Networks of Models: Analogy as the Engine Isomorphism-Talk Conceals

Sections 11.2 through 11.8 have tested the map-as-model analogy against several accounts, from da Costa and French's partial structures to Morgan and Morrison's models as mediators, each asking how closely a scientific model's relation to its target must resemble a map's relation to the terrain it charts. Mary Hesse's Models and Analogies in Science, predating most of the accounts surveyed in this chapter, approaches the question from underneath the isomorphism vocabulary those accounts share, distinguishing within any model-target comparison a positive analogy, the properties model and target are known to share; a negative analogy, the properties known to differ; and, doing the actual explanatory and predictive work, a neutral analogy, the properties not yet known to be shared or not, whose exploration is precisely how a model generates novel, testable predictions rather than merely restating what was already known about the target.

The cartographic case makes the three-way distinction unusually concrete, because a map's positive, negative, and neutral analogies with its terrain are seldom left implicit the way a scientific model's often are. A road map's positive analogy with the terrain is connectivity and relative sequence, preserved by construction; its negative analogy, flagged by the very existence of a legend and a stated projection, is metric accuracy, known in advance to be sacrificed along the lines Chapter 4 proved must be sacrificed; its neutral analogy, smaller and more interesting than either, is whatever the map's user has not yet checked against

the terrain but might reasonably infer from it, a shortcut inferred from apparent proximity on the page that turns out, or fails to turn out, to hold on the ground. Section 11.5's question, what actually transfers from Tissot's theorem and what does not, is answered by Hesse's scheme with a precision the isomorphism-versus-partial-isomorphism debate alone does not supply: what transfers is the demonstrated existence of a nonempty negative analogy in any sufficiently rich model, cartographic or scientific, and what does not transfer is any particular content of that negative analogy, which is domain-specific in exactly the way Chapter 4's geometric proof is and a general theory of modeling cannot be.

Hesse's account also supplies a mechanism this chapter's more static isomorphism-talk has lacked throughout: the neutral analogy is not a residual category to be tolerated but the site where a model does its distinctive work, generating hypotheses about the target precisely where the comparison has not yet been settled either way. A cartographer inferring an unsurveyed river's likely course from the drainage pattern legible elsewhere on the same map is exploiting a neutral analogy in Hesse's exact sense, and the inference's eventual vindication or defeat, once the river is actually surveyed, moves that piece of the comparison from neutral into positive or negative analogy without touching the rest of the map's validity. Read this way, the map-as-model analogy this chapter has been testing throughout is at its strongest not where model and target are shown to agree, nor merely where they are shown to differ, but at Hesse's neutral middle, which is where Wimsatt's robustness analysis in Section 11.7 and Morgan and Morrison's mediating models in Section 11.8 both turn out, on inspection, to have been operating all along.

11.9 Kuhn's Paradigms and the Cartographic Model's Historical Discontinuities

Chapter 11 has so far treated the relation between cartographic and scientific models mostly synchronically, comparing frameworks (partial structures, similarity views, constructive empiricism) as they stand rather than as they

succeeded one another historically. Thomas Kuhn's account of paradigm shifts in The Structure of Scientific Revolutions supplies the missing diachronic dimension, and cartography's own history supplies an unusually clean test case for it.

The transition from a Ptolemaic conic-projection tradition (Section 1.8) through Mercator's conformal projection to the diverse, trade-off-aware pluralism Chapter 4 describes has some features of a Kuhnian paradigm shift — Mercator's projection was adopted for reasons (navigational utility, rhumb lines as straight lines) that had little to do with the criteria Section 4.6's Pareto frontier would later make explicit, and its later criticism (Section 13.7's Peters controversy) was fought partly as a battle over which virtues even counted as relevant, a hallmark of what Kuhn calls incommensurability between successive frameworks' standards of evaluation.

The disanalogy is instructive and worth stating with the same care Section 11.5 already brought to what transfers from Tissot's theorem and what does not. Kuhn's paradigm shifts are typically read as involving some genuine loss alongside gain, evaluative standards that do not straightforwardly nest inside one another; cartography's history, once Chapter 4's theorem is available, does not work this way, since Tissot's result is not itself paradigm-relative and supplies a framework-neutral map of the entire space of trade-offs within which Ptolemy, Mercator, and every later projection can be located and compared on common terms. Cartography's history therefore looks Kuhnian in its sociology of adoption and controversy, Section 13.7 included, while remaining decidedly cumulative, not incommensurable, at the level of the mathematics beneath it — a distinction this volume's companion series makes it unusually easy to draw.

11.10 Cartwright's Dappled World: Patchwork Realism Against a Single Global Structure

Section 11.6 traced what survives of structural realism once Tissot's theorem is taken seriously as a limiting case. A further and in some ways more radical

position from the same debate, Nancy Cartwright's patchwork or 'dappled world' view, denies not merely that any single representation can capture nature's full structure, which Chapter 4 already establishes for the cartographic case, but that nature itself has a single unified structure of laws waiting to be captured at all, proposing instead that the world is composed of numerous domains, each governed by locally reliable regularities that need not fit together into one global system.

The cartographic analogy this volume has built across Part II supplies a strikingly literal illustration of what Cartwright otherwise argues for by more abstract means. Section 5.4's discussion of local accommodations to the hairy ball theorem — Section 7.7's cognitive maps, Section 7.8's etak, Section 7.10's echolocation — are, precisely, patches: locally reliable, locally sufficient representational strategies that never claim to compose into one global field, because Chapter 5 proves no such field exists to compose into. Cartwright's dappled world can be read, cautiously, as generalizing this volume's most concrete finding about spheres and vector fields into a broader metaphysical picture of scientific representation generally.

The caution is necessary and the Preliminary Note's discipline requires stating it plainly: Chapter 5's conclusion is a theorem about a specific mathematical object, the sphere, and Cartwright's is a much more general and considerably more contested claim about the structure, or lack of it, of nature as a whole. What this volume's cartographic case can offer Cartwright's debate is not evidence for the metaphysical thesis, which would overstate what a topological theorem about map-making can show, but a fully worked, non-question-begging existence proof that patchwork representation without global unification is sometimes not a regrettable limitation but, as Section 5.4 already argued, the mathematically forced and only available option.

11.11 A Case in Point: Climate Model Ensembles and the Institutionalized Refusal of a Single Global Model

The philosophical debate Sections 11.7 and 11.10 have just traced, Wimsatt's robustness analysis and Cartwright's dappled world, has a fully institutionalized contemporary instance in the practice of climate modeling, and specifically in the Coupled Model Intercomparison Project (CMIP), the coordinated framework under which dozens of independently built general circulation models, from centers including the U.S. National Center for Atmospheric Research, the UK Met Office's Hadley Centre, the Max Planck Institute for Meteorology, and others, are run on common experimental protocols and compared rather than reconciled into one another. Each model encodes its own choices about cloud parameterization, ocean mixing, aerosol forcing, and dozens of other subgrid processes that cannot be resolved from first principles at planetary scale, choices that differ across modeling centers for reasons of scientific judgment, institutional history, and available computational resources rather than because any one center has identified an error the others have not yet corrected. The Intergovernmental Panel on Climate Change's assessment reports do not resolve this plurality by selecting a single best-performing model; they report the ensemble's spread as itself the scientifically meaningful quantity, treating agreement across structurally different models as evidence of a robust finding and disagreement as a flagged uncertainty, in a procedure that is Wimsatt's robustness analysis, Section 11.7 already described in the abstract, executed as standing international infrastructure rather than as a philosopher's illustration.

The disanalogy is worth stating with the same care Section 11.10 already brought to Cartwright's own case, because the fit is close but not exact. Cartwright's patches, and the cartographic patches of Sections 5.4, 7.7, and 7.8 that Section 11.10 read alongside them, are typically non-overlapping: each covers a distinct local domain, a distinct region of the sphere or a distinct wayfinding context, and no patch claims jurisdiction over another's territory. The CMIP ensemble is not that: every model in it claims the same domain, the entire global climate system, and their divergence is not a partition of coverage but a genuine disagreement about the same object, retained deliberately rather than resolved. What the climate case actually illustrates, then, is closer to Section 11.7's robustness thesis

than to Section 11.10's metaphysical patchwork: not that the world lacks a single unified structure, which the ensemble's convergence on well-constrained variables like global mean temperature trend would seem to count against, but that triangulating a contested estimate through several independently flawed models, none individually trusted, can be more reliable than any single model's claim to have gotten the whole picture right, which is a methodological lesson about representation under irreducible uncertainty rather than a claim about nature's own composition.

Chapter 12 — Explanatory Models and the Trade-off Between Accuracy and Understanding

12.1 A second trade-off, not the same trade-off

Chapter 11 searched the philosophy of scientific representation for a second Tissot-style theorem and found, instead, a family of independently-motivated accounts sharing cartography's three-part shape without sharing its proof. A related but distinct literature raises a different trade-off, one this volume has not yet examined directly: not accuracy against accuracy, as angle trades against area within Tissot's theorem, but accuracy against explanatory power, the finding, defended across several independent philosophical projects, that a model can become a worse explanation precisely by becoming more accurate, because added realistic detail can bury the pattern that actually does the explanatory work under detail that does not. This chapter examines that trade-off with the same discipline Chapter 11 applied: is it, anywhere, provable and purpose-independent in Tissot's sense, or is it real but relative to what a finite mind can take in — a fact about explanation for reasoners like us, rather than a fact about the target system the way curvature is a fact about the sphere? Stated in advance: one account examined here, Batterman's, comes closer than anything in Chapter 11 to a Tissot-like asymmetry, and Section 12.2 explains carefully why the resemblance is real without being the same kind of necessity.

12.2 Minimal models and Batterman's asymptotic explanations

Batterman's work on minimal models examines cases where a strikingly simple, deliberately unrealistic model explains a pattern shared by physically dissimilar systems — the ubiquity of the Gaussian distribution across independent sources of noise, or the identical critical-point behavior of fluids and magnets near a phase transition, despite their having almost nothing else in common physically. Batterman's central claim is not merely that the simple model is easier to use than a detailed one, but that it explains better, because the mathematics of the relevant asymptotic limit is itself what is responsible for the universal pattern, and a more detailed, more realistic model actively obscures this by reintroducing system-specific detail that the asymptotic limit washes out. On this view, realism is not merely unnecessary for this kind of explanation; it is an obstacle to it.

This is the closest approach, anywhere in this chapter or the last, to a Tissot-shaped asymmetry: a case where more of one virtue (descriptive realism) provably costs a different virtue (explanatory transparency), rather than the two simply varying independently. The comparison must nonetheless be marked carefully rather than affirmed too quickly. Tissot's impossibility is a fact about the target — the sphere's curvature — that holds regardless of any observer, explanatory aim, or cognitive limitation; a civilization that thought in ways utterly unlike human cognition would still find no isometric projection of a sphere onto a plane. Batterman's impossibility is a fact about what makes a pattern explanatorily transparent to a finite reasoner working with asymptotic mathematics; it is a claim about explanation, an epistemic and pragmatic notion, not a claim built into the target system the way curvature is built into the sphere independent of anyone's cognition. The asymmetry is real and, within its domain, close to lawlike; its source is not geometric, and the disanalogy is worth holding onto exactly because the resemblance is otherwise so tempting to overstate.

12.3 Elgin's felicitous falsehoods and Potochnik's idealization for understanding

Catherine Elgin's notion of a felicitous falsehood names a model that is strictly false of its target — a frictionless plane, an ideal gas, a perfectly rational agent — and yet affords genuine understanding precisely by departing from literal truth in a controlled, acknowledged way, exemplifying features that a fully accurate description would leave harder to see. Angela Potochnik develops a related and more systematic thesis: idealization is not a regrettable concession forced on science by finite time and computing power, to be eliminated as soon as resources allow, but constitutive of what causal and explanatory understanding actually requires, because understanding depends on isolating a difference-making pattern from a welter of detail that is causally present but explanatorily beside the point. On this view, a maximally accurate model, even a hypothetically achievable one, would often be a worse vehicle for understanding, not merely an unneeded luxury once understanding is already secured by other means.

This volume anticipated a version of this point from within cartography specifically, well before this chapter: Section 2.2's discussion of Borges's fable of the map at the scale of the territory, and of Carroll's earlier version of the same joke, already showed that a map preserving everything would be useless as a map, its completeness itself the defect. Elgin and Potochnik generalize the identical insight for the epistemology of science broadly, arriving at it independently and for reasons proper to their own domain, which strengthens — without proving — the analogy in exactly the way Section 11.5 described: independent convergence on the same shape, from writers not discussing cartography, counts as evidence that the shape is a real feature of representation generally, even where, as here, no second Tissot theorem underwrites it.

12.4 Cartwright's simulacrum account and the honesty of a governing fiction

Section 10.6 introduced Cartwright's nomological machines and her simulacrum account of models as engineered, idealizing constructions built to let a specific law operate cleanly, and noted there that Cartwright presses the view into the

strong claim that most physical laws, read literally, are true only of the models built to house them and not of the world beyond those models. The same caution registered there applies with equal force here, where the topic is explanation rather than idealization narrowly construed: Cartwright's simulacrum account denies that models aim at truth at all, treating explanatory success and literal truth as coming apart in principle, not merely in current practice. This is a considerably stronger and more contested thesis than anything Elgin or Potochnik requires, and stronger than anything this volume's own treatment of cartography has argued or needs to argue.

The distinction is worth drawing sharply because it is exactly the kind of overreach Section 1.2 warned against under the name of the inflationary temptation. Nothing in Tissot's theorem or the hairy ball theorem requires denying that a projection is honestly evaluable as more or less accurate in Section 3.2's third, purpose-relative sense, and nothing in them requires denying that the geographic facts a map depicts hold true independent of any map at all; cartography's impossibility theorems constrain representation, they do not dissolve the reality being represented. Cartwright's fictionalism about physical law is a further, separable philosophical commitment, and this chapter borrows her structural account of purpose-built idealization — which resonates genuinely with Chapter 6's conventionalism and Section 4.3's choice among survivors — without adopting the wider anti-realism, flagging the boundary explicitly rather than letting a useful borrowed vocabulary carry an unargued conclusion along with it.

12.5 Reading the two trade-offs side by side without collapsing them

Two trade-offs have now been examined at length and should not be allowed to blur into one. The first, Tissot's, is the subject of Chapter 4: a fact about curvature, exact, purpose-independent, provable, holding for any projection whatsoever regardless of who uses it or why. The second, surveyed across this chapter, is a trade-off between descriptive accuracy and explanatory or epistemic understanding, real and independently attested across several distinct

philosophical projects — Batterman's, Elgin's, Potochnik's, Cartwright's — but resting, in every case examined, on facts about explanation for finite reasoners rather than on facts about the target system the way Tissot's theorem rests on the sphere's intrinsic curvature. Batterman's minimal models came closest to a lawlike asymmetry, and even there the necessity traced back to asymptotic mathematics doing explanatory work for a finite mind, not to a geometric fact holding independent of any mind at all.

Holding both trade-offs up together, without collapsing the second into the first, is exactly the exercise Chapter 3 opened this volume's engagement with representation by calling for, and Chapter 11 practiced at length: using a hard-won, fully proved result to illuminate a genuinely similar-shaped but unprovable philosophical claim, while marking, at every step, which register — rigorous constraint or careful analogy — is doing the work. That two such different trade-offs, arrived at by entirely different routes, should exhibit so similar a shape is itself, as Section 11.5 already argued, a finding worth taking seriously. It is not, and this volume has tried throughout not to let it become, a license to treat the second trade-off as though Chapter 4 had already proved it too.

12.6 Bounded Rationality: Simon's Version of the Same Trade-off

Section 12.2 found, in Batterman's minimal models, the closest approach anywhere in this chapter to a Tissot-shaped necessity binding accuracy against understanding, and traced its source to asymptotic mathematics rather than to any fact about a target system's intrinsic geometry. A second, more general source for the same trade-off is worth adding, owed to Herbert Simon's account of bounded rationality, because it locates the necessity in a place structurally different again from Batterman's asymptotic limits: not in the mathematics of a particular explanatory pattern, but in the finite computational and memory resources of any reasoner, however that reasoner is otherwise constituted. Simon's central claim, developed across Administrative Behavior and his later work on bounded rationality, is that any real decision-maker or model-builder

operates not by finding the objectively optimal representation of a problem but by satisficing — constructing a representation simple enough to reason with inside the reasoner's actual computational limits, then choosing an option that meets some threshold of adequacy rather than searching exhaustively for the true optimum, because the exhaustive search a fully accurate representation of the problem would require is not available to any reasoner with finite time and finite working memory, human or artificial.

Read alongside Section 12.2, Simon's account supplies the general theory of which Batterman's minimal models are a specific instance: Batterman's asymptotic explanations are one particular way a representation gets simplified in order to become tractable for a finite mind, but Simon's claim is broader, holding that some such simplification is forced wherever a bounded reasoner represents a problem at all, independent of whether the specific problem happens to admit an asymptotic limit doing Batterman's kind of explanatory work. This generality is also exactly where the disanalogy with Tissot's theorem becomes sharpest, more so than for Batterman's narrower case: Simon's bound is a fact about the reasoner, not about the target system in any sense at all, so that a hypothetical reasoner with unbounded computational resources would face no version of Simon's trade-off, however complex the underlying system, whereas Tissot's theorem would bind such a reasoner exactly as it binds a human cartographer, since the sphere's curvature is indifferent to how much computation is brought to bear on it. Simon's trade-off is, in this precise sense, the least geometric and most purely epistemic of every case Chapters 10 through 12 have examined — real, well-attested across decision theory and cognitive science, and entirely a fact about reasoners rather than a fact about the world the reasoner is trying to represent, which is exactly why it belongs at the far end of the spectrum this chapter's closing section, 12.5, already asked the reader to keep in view.

12.7 Strevens' Kairetic Account: Difference-Making as a Third Trade-off

Sections 12.2 through 12.6 have examined several accounts of why a simplified, partly false model can nonetheless explain: Batterman's asymptotic explanations, Elgin's felicitous falsehoods, Potochnik's idealization for understanding, Cartwright's simulacrum account, and Simon's bounded rationality. Michael Strevens' kairetic account, developed at length in Depth: An Account of Scientific Explanation, supplies a further, more procedurally explicit criterion worth adding: an explanation is improved, on Strevens' account, by removing any causal detail whose removal does not change whether the outcome to be explained still occurs, a systematic stripping-away he calls the kairetic criterion, continued until only the causal difference-makers, the details the outcome actually depends on, remain. A model built this way is not simplified for convenience or tractability, in Simon's sense, nor idealized toward a limiting case, in Batterman's; it is deliberately pared down to exactly the causal structure doing the explanatory work, with everything else, however true, treated as explanatorily inert and properly omitted.

Cartography supplies the kairetic criterion an unusually direct test case, because a map's omissions are already, independently of this literature, organized by exactly this logic. A subway map's omission of true geographic distance is not a concession to drafting convenience, in Simon's register, nor a limiting idealization, in Batterman's; for a rider deciding which line to take, true distance is, in Strevens' precise sense, not a difference-maker — changing it, within the range any real city's geography actually presents, would not change the correct answer to the rider's question, which depends only on connectivity and sequence. A nautical chart's comparatively fastidious preservation of bearing and distance, by contrast, is exactly what Strevens' criterion would recommend keeping for a navigator, since an error in either is precisely the kind of detail a collision or a grounding does depend on. What Section 4.3 described as purpose disciplining the choice among Tissot's survivors, and what Section 10.3 called the difference between a flaw and a cost, both restate, in cartography's own vocabulary, the same criterion Strevens formalizes generally: keep what the outcome the map is for actually depends on, and the kairetic account supplies

the missing general name for why that selection, rather than any other, counts as improving rather than merely shortening the representation.

The trade-off this adds is a third, and distinct from the two Section 12.1 was careful to keep apart from each other. Chapter 4's trade-off is forced by geometry and admits no purpose-independent optimum; Section 12.1's accuracy-understanding trade-off is forced by the cognitive cost of tracking every causal detail at once. The kairetic trade-off is a third thing again: not a cost of tracking too much, but a risk of tracking too little, since an aggressive kairetic stripping-away can remove a detail later discovered to be a difference-maker after all, exactly the failure mode Section 10.4 warned distortion-as-condition must not be allowed to license. Strevens' own discipline on this point matches this volume's throughout: the kairetic criterion is answerable to the actual causal structure of the target system, not to the modeler's antecedent sense of what seems irrelevant, which is why a subway map's omissions are vindicated by the actual irrelevance of distance to route choice, and would not be vindicated merely by a mapmaker's hunch that riders wouldn't care.

12.8 Marr's Levels of Analysis: A Third Register for the Accuracy–Understanding Trade-off

Chapter 12 developed its central trade-off, between accuracy and understanding, through Batterman's asymptotic explanations, Elgin and Potochnik on felicitous idealization, and Cartwright's simulacrum account. David Marr's tri-level framework for cognitive science, developed independently in vision research, offers a further vocabulary for the same trade-off, distinguishing the computational level (what problem is being solved, and why), the algorithmic level (what representation and procedure solves it), and the implementational level (how that procedure is physically realized).

Marr's own methodological point, that a full understanding of a system requires all three levels but that most explanatory progress in vision research came from getting the computational level right before worrying about implementation,

maps cleanly onto Section 12.2's treatment of minimal models: a minimal model, on Marr's terms, is often pitched deliberately at the computational level, explaining why a trade-off must hold without committing to any particular algorithmic or implementational story, exactly the pattern Section 12.2 identified in Batterman's asymptotic explanations of universal phenomena.

The connection back to this volume's cartographic core is direct rather than merely illustrative. Chapter 4's theorem operates at Marr's computational level: it explains why any projection must trade conformality against equal-area representation, independent of which specific algorithm generates the projection's coordinates or which specific software implements that algorithm on a given machine. Section 12.5's instruction to read Chapters 4 and 12's two trade-offs side by side without collapsing them gains a further tool from Marr: the two trade-offs occupy the same level of analysis, the computational one, which is precisely why they resemble each other structurally while remaining, as Section 12.5 insists, genuinely distinct trade-offs rather than one trade-off wearing two names.

12.9 Kolmogorov Complexity and a Formal Measure of the Simplicity Chapter 12 Has Assumed

Chapter 12's trade-off between accuracy and understanding has relied throughout on an intuitive sense of simplicity — Section 12.2's minimal models, Section 12.3's felicitous falsehoods, Section 12.6's bounded rationality all invoke a model or heuristic being simpler than some more accurate alternative without the chapter committing to any formal measure of what simplicity is. Kolmogorov complexity supplies one candidate: the complexity of a string (or, by extension, a model's description) is the length of the shortest program that could produce it, giving simplicity a precise, if generally uncomputable, meaning.

This connects directly to a formal cousin of Section 10.8's Akaike information criterion, minimum description length, which recasts model selection as a search for the model minimizing the combined cost of describing the model itself plus

describing the data's remaining deviations from it — an explicit trade-off between the two lengths that gives Section 12.5's instruction to read Chapters 4 and 12's trade-offs side by side a further formal echo: just as Tissot's theorem trades conformality against equal-area representation along a provable frontier, minimum description length trades model simplicity against residual error along a frontier that is, in principle, exactly computable for any given coding scheme, even though Kolmogorov complexity itself, in general, provably is not.

The uncomputability is worth stating rather than glossing over, since it marks a genuine disanalogy with Chapter 4's case. Tissot's theorem gives an exact, checkable answer for any named projection; Kolmogorov complexity is formally uncomputable in general, meaning no algorithm can determine the shortest possible description of an arbitrary object, so the measure that makes Occam's razor mathematically precise is, in its full generality, one that cannot actually be applied. What Chapter 12's working notion of simplicity gains from the comparison is not a usable formula but a clarified sense of its own limits: the intuitive judgments of simplicity operating throughout this chapter are approximations to something formally exact and formally out of reach, not merely to something vague.

12.10 No Free Lunch Theorems: A Genuinely Proven Cousin, Outside Explanation Altogether

Section 12.1 opened this chapter by asking whether the trade-off between accuracy and understanding is anywhere provable and purpose-independent in Tissot's sense, or only real and relative to a finite reasoner, and Sections 12.2 through 12.9 answered, case after case, that it is the latter — Batterman's asymptotic explanations came closest, but even there the necessity traced back to mathematics doing explanatory work for a mind, not to a fact about the target system the way curvature is a fact about the sphere. One further case, from computer science rather than the philosophy of explanation, is worth closing the chapter with precisely because it breaks this pattern: Wolpert and Macready's No Free Lunch theorems for optimization, proved in 1997, state that averaged over

all possible objective functions, every search or optimization algorithm performs identically — no algorithm can outperform random search when its performance is averaged across the entire space of problems it might be asked to solve, so that any algorithm's superiority on some class of problems is necessarily balanced by inferiority on another class, exactly compensating.

This is not, on inspection, a fact about finite reasoners in Simon's sense (Section 12.6) or about asymptotic mathematics in Batterman's (Section 12.2); it is a theorem about the structure of the space of all possible problems itself, holding for any algorithm whatsoever, including a hypothetical one with unbounded computational resources, in exactly the way Section 12.6 noted Tissot's theorem, but nothing else examined in this chapter, would still bind a reasoner of unlimited power. An algorithm that seems to 'understand' a problem better than blind search, in the sense of solving instances of it faster or more reliably, is not thereby a better algorithm in any purpose-independent sense; it is an algorithm whose built-in assumptions happen to match the structure of the particular problems it is actually asked to solve, at a cost, guaranteed by the theorem, of performing worse on problems lacking that structure. This is Section 4.3's lesson about choosing among Tissot's survivors, restated in a domain with no geometry in it at all: there is no algorithm optimized for solving problems in general, only algorithms optimized for the specific, structured subset of problems a designer actually cares about, exactly as there is no projection optimized for mapping in general, only projections optimized for a stated purpose.

What this closing case adds to Section 12.5's instruction to hold the volume's two trade-offs apart without collapsing them is a third, cleanly separated data point rather than a blurring of the first two. The No Free Lunch theorems are exact, proven, and purpose-independent in exactly Tissot's sense, yet they concern neither cartography nor explanation but the combinatorics of search — evidence, alongside Section 5.8's Arrow and Section 5.10's Brouwer, that a genuinely Tissot-shaped result can turn up almost anywhere once the space of a problem is specified precisely enough, and a useful corrective to any impression the rest of this chapter might have left that provable, purpose-independent trade-offs are

cartography’s peculiar property rather than a shape mathematics is simply disposed, now and again, to produce.

12.11 A Closing Case: Interpretable and Black-Box Models as the Trade-off Made an Engineering Choice

Sections 12.2 through 12.10 have approached the accuracy–understanding trade-off mostly through cases the discipline of philosophy of science had already assembled, minimal models, simulacrum accounts, bounded rationality, Marr's levels. A case from outside that literature, close to contemporary practice rather than to its philosophical reconstruction, states the same trade-off as an engineering choice made daily rather than an interpretive puzzle argued over after the fact. Contemporary machine learning offers, on one side, models, large ensembles of decision trees, deep neural networks with many layers, whose predictive accuracy on held-out data routinely exceeds that of simpler alternatives, and on the other, models, linear regressions, shallow decision trees, rule lists, whose internal structure a human reader can trace from input to output without a further, separate explanatory apparatus. The field's own vocabulary for this, the accuracy-interpretability trade-off, names exactly the tension Section 12.1 distinguished from Chapter 4's Tissot trade-off: not a cost paid in the model's fit to the world, but a cost paid in a human being's capacity to hold the model's reasoning in view, precisely Section 12.2's distinction between a model that gets the answer right and a model that lets a reader see why.

Marr's levels, Section 12.8's framework, apply to this case with unusual directness: a black-box classifier can be fully specified at the computational and algorithmic levels, its objective and its procedure both exactly stated, while remaining opaque at the level a working analyst actually needs, an implementation a human can inspect and predict the behavior of without running it. The interpretable-model literature's own preferred remedies, post-hoc explanation methods that approximate a complex model's local behavior with a simpler, human-readable one, are a direct engineering analogue of Elgin's felicitous falsehood and Potochnik's idealization for understanding, Section 12.3's

cases: a known, acknowledged simplification substituted for the true generating process because the true process, though more accurate, cannot be held in a human mind at all. Whether the field's growing preference, in regulated domains such as medicine and lending, for interpretable models even at some cost in accuracy amounts to a considered philosophical judgment about which side of Chapter 12's trade-off matters more, or merely a legal and institutional response to accountability requirements, is a question this volume is not positioned to settle; what belongs here is only the observation that a trade-off this chapter has argued for on largely a priori grounds is now also a line item in machine learning conference proceedings, argued over with the same structure and none of the philosophical vocabulary.

Chapter 13 — Cartography and Power: The Map as Political Act (with Historiographical Caution)

13.1 Two very different reasons a map can be wrong

A map can fail to show the territory accurately for two entirely different reasons, and the literature on maps and power has not always kept them apart. The first reason is the one this volume has spent three parts establishing: no map, however carefully made, can avoid some distortion, because the distortion is a mathematical consequence of representing a curved surface on a flat one, or an orientation field on a surface with the wrong topology. The second reason is not mathematical at all: a map can distort in a particular direction, emphasize a particular region, omit a particular feature, or choose a particular center, because the person or institution making it wanted it to look that way for reasons that have nothing to do with Gauss's Theorema Egregium. The first kind of distortion could not have been avoided by anyone, however well-intentioned; the second kind could have been, and was not.

13.2 What critical cartography got right, and its own overreach

The scholarship associated with critical cartography, J.B. Harley's essays being the usual point of entry, is right about a great deal: maps have routinely served colonial administration, military planning, and territorial claims, and the appearance of neutrality — the map as a simple picture of what is there — has itself been a rhetorical resource for interests that were anything but neutral. The Mercator projection's disproportionate enlargement of high-latitude, historically colonizing regions relative to the equatorial regions it colonized is frequently cited in this literature as an instance of cartographic power. This volume does not dispute that Mercator's projection enlarges high latitudes — that is a direct, quantifiable consequence of the conformal requirement Chapter 4 already derives from Tissot's theorem — nor that the choice to use it for a general-purpose world map, long after its actual purpose of constant-bearing navigation had stopped requiring it, reflects a further choice with its own history and its own consequences worth criticizing. What this volume resists is a slide, present in some of the more polemical versions of this literature, from 'this particular choice was political' to 'the distortion itself was the politics' — a move that quietly treats a Tissot-mandated cost as though it were itself the colonial act, when the colonial act was a separable choice made within the space Chapter 4 shows any such choice must occupy. The mathematics did not choose Mercator for the atlas publishers; it only guaranteed that whatever they chose would distort somewhere.

13.3 The connection to Chapter 6: choice within genuine constraint

Chapter 6 already gave the general form this chapter now applies to a harder case. Poincaré's conventionalism showed that a choice constrained by genuine facts — there, the empirical underdetermination of geometry; here, the impossibility results of Part II — is not thereby an arbitrary choice, free of all further explanation. It still needs one, and the explanation is where the interesting work is. Exactly the same structure holds for the choice among projections, among what a map centers, and among what a map omits. The impossibility theorems fix the menu; they do not fix the order. Recognizing that a

chef did not invent the rule that some dishes cannot be both very sweet and very sour at once does not excuse the chef from having chosen this dish, for this table, for these reasons — and a diner who wanted a different balance is entitled to ask why this one was served instead. The genuinely proven part of the situation and the genuinely contingent part of the situation are both real, and treating either as though it dissolved the other is the specific error this chapter is written to head off, in both directions: the mathematics is not an alibi for the choice, and the choice's motives are not a refutation of the mathematics.

13.4 Power exploits the residual freedom, not the impossibility

The precise, and more defensible, claim this volume can make on behalf of the critical-cartography literature is therefore narrower than 'maps distort because they serve power', but not, for that reason, weaker: power operates in the residual freedom the impossibility theorems leave open, and it operates there more effectively for being able to point to that freedom's necessity as though it excused its use. Every map must choose a projection, a center, a scale, an inclusion, an omission; nothing in the mathematics says which. That 'must choose' is where a cartographer's interests, a sponsoring state's interests, or an unexamined convention gets to act, and the fact that some choice was unavoidable does the political chooser the considerable favor of making the particular choice look unavoidable too. Seeing this clearly is, if anything, a stronger position from which to criticize a given map's politics than the blurred version, because it locates the responsibility precisely: not in the theorem, which explains why any map distorts, but in the choice, which explains why this map distorts this way, and asks — as Chapter 6 already insisted a constrained choice must always be asked — for its reasons.

13.5 A Counter-Example: When a Choice Is Not Political

Sections 11.2 through 11.4 have been occupied with distinguishing genuine political choice from mathematically mandated distortion, and with insisting that

the two not be collapsed into one another. That insistence would be one-sided if it only ever ran in the direction of finding hidden politics; the same discipline requires being equally willing to find its absence, and a brief counter-example is worth setting alongside Section 13.2's Mercator case for exactly that reason. The Web Mercator projection, used by nearly every major online mapping service for its base tile layer, is mathematically a variant of the historical Mercator projection Section 13.2 already discussed, and inherits the same high-latitude area inflation; a critic working only from Section 13.2's framework might be tempted to read its near-universal adoption by twenty-first-century technology companies as a further instance of the same colonial-cartographic pattern, updated for the internet age. The historical record of why Web Mercator was chosen does not support that reading nearly as well as it supports a different, far less politically loaded explanation: square map tiles, the standard unit of web-based slippy maps, need a projection under which north stays reliably 'up' and a rectangular tile at any zoom level stays a rectangle, and conformal projections are exactly the family with that property, because conformality is precisely the preservation of local angle and hence of shape at a point. Mercator's specific conformal projection, moreover, has a simple closed-form relationship between latitude, longitude, and pixel coordinates that made it computationally cheap on the limited hardware early web-mapping systems ran on, an engineering constraint with no discernible connection to the nineteenth-century atlas politics Section 13.2 examined, even though the resulting map inherits the identical area distortion.

This does not undo Section 13.4's claim that power operates in the residual freedom the impossibility theorems leave open; it refines it, by showing that the same residual freedom can be filled by considerations that are not political at all, and that telling the two apart requires exactly the kind of specific historical inquiry into why a particular choice was made that this volume has tried to model, rather than a general suspicion that any distortion, once traced to a choice, must trace further to an interest. A reader equipped only with Section 13.2's Mercator case risks over-applying its lesson to every subsequent instance

of the same projection; Web Mercator's very different history is offered here as a corrective, and as a reminder that the chapter's argument was never that all cartographic choices are political, only that some are, and that the impossibility theorems alone cannot tell a reader which.

13.6 Cartograms: When Distortion Is the Point, Declared as Such

Section 13.5 offered a counter-example to caution against over-applying Section 13.2's Mercator case; a complementary case, at the opposite extreme, is worth adding to close this chapter's treatment of the residual freedom Section 13.4 identified as power's actual point of purchase. A cartogram — a map in which the area or shape of each region is deliberately resized in proportion to some non-geographic variable, population, GDP, votes cast, carbon emissions — is, judged by ordinary geographic accuracy, about as distorted as a map can be: a cartogram of world population inflates India and China far beyond their true land area and shrinks Australia and Canada far below theirs, a deformation more extreme than anything Mercator's conformal requirement produces. And yet a cartogram is, in an important sense Section 13.1's opening distinction was built to capture, the least deceptive kind of map this volume has examined, because the distortion is not merely present but is the entire declared content of the representation: no reader consulting a population cartogram is given any reason to believe it preserves true area, since preserving true area is not the property the map claims, states in its legend, or was constructed to have.

This is the sharpest available illustration of Section 13.4's claim that power operates in the residual freedom the impossibility theorems leave open, not in the impossibility itself, because a cartogram shows the same residual freedom used for the opposite of concealment: Mercator's high-latitude inflation is a side effect of a projection chosen for a different, stated purpose, constant-bearing navigation, and only later criticized for continuing to structure a general-purpose atlas long after that purpose stopped applying, whereas a cartogram's area distortion is not a side effect of anything; it is manufactured on purpose, in full view, as the only way to make a non-geographic quantity legible on a map at all,

since Chapter 3's homomorphism already established that some property of the territory must be sacrificed for any other to be shown clearly. A map's distortion being large, in other words, tracks nothing at all about whether that distortion is being used honestly or dishonestly; a cartogram is enormously distorted and maximally honest about it, while a conventional-looking equal-area map can misinform simply by choosing, without disclosure, a center or a set of omissions that a critical reading only uncovers with effort. What decides the political question Section 13.1 opened this chapter with is never the mere fact or magnitude of distortion, which every map must have somewhere, but whether the specific distortion chosen is disclosed, purpose-fitted, and available for the reader to evaluate — the same standard, transposed once more, that Section 6.3 already applied to the choice of projection in general, and the standard this chapter leaves the reader with as it closes.

13.7 The Peters Controversy: A Case Where the Critique Misused the Very Theorem It Invoked

Sections 13.1 through 13.6 have argued a specific, disciplined thesis: that political critique of cartography is correct to notice that choices are made within the space Chapter 4's theorem leaves open, and mistaken when it treats that residual freedom as evidence the underlying constraint is itself an illusion or a pretext. The controversy Arno Peters provoked from the mid-1970s onward, when he presented what he called a new, more equitable equal-area projection and argued that Mercator's enormous popularity had for centuries visually inflated the high-latitude, historically colonizing nations at the expense of the equatorial, historically colonized ones, is worth treating as a full case study of this thesis, precisely because it contains a defensible core and a technically confused overreach, and because a number of aid and educational organizations that adopted the Peters projection in the 1980s did so on the strength of the overreach as much as the core.

The defensible core survives Chapter 4's scrutiny intact and restates Section 13.4's argument exactly: Mercator's near-universal adoption as the default world

map for non-navigational purposes, classroom wall maps prominent among them, was indeed a choice, not a neutral default forced by the mathematics, and the choice does have the consequence Peters described, a systematic visual inflation of high-latitude landmasses relative to equatorial ones, Greenland famously appearing comparable in size to Africa though it is in reality roughly one-fourteenth its area. Nothing in Chapter 4's theorem required this particular point on the trade-off frontier to become the classroom default; Section 6.4's conventionalism already supplies the resources to say that a different, equally defensible point could have become standard instead, and that the fact it did not is a historical and institutional fact worth the critical attention Peters gave it.

The overreach is equally instructive and is exactly the failure Section 13.2 anticipated in general terms. Peters presented his own construction as though it were uniquely accurate, a corrective to Mercator's distortion rather than one further point on the same Pareto frontier Section 4.6 described; professional cartographers' objections, sometimes dismissed at the time as a defense of the establishment's preferred map, were substantially technical and correct: Peters' projection, an equal-area cylindrical construction, purchases its equal-area property at a cost in shape distortion, elongating and compressing landmasses near the equator and the poles in ways many working cartographers judged more visually distorting than Mercator's own concession, not less distorted in some purpose-independent sense Chapter 4's theorem shows does not exist. Both maps are exactly and only what Chapter 4 predicts any projection must be, correct along one axis at the necessary cost of another; the controversy's real lesson, on this volume's terms, is not that Peters was right and Mercator wrong, or the reverse, but that a genuine political critique of Mercator's historical dominance was, in the event, bundled with a technically false claim to have escaped Tissot's theorem altogether, and the bundling did the critique's legitimate half no favors once the technical half was shown, correctly, to be confused.

13.8 Cartographic Silence: The Ethics of What a Map Leaves Out

Chapter 13 followed Harley's reading of the map as a political act largely through what maps assert — the boundaries drawn, the names printed, the symbols chosen to carry a particular framing. A second and related category, cartographic silence, what a map deliberately or systematically fails to show, deserves separate treatment, since Section 13.4's claim that power exploits the residual freedom Chapter 6 leaves open applies at least as forcefully to omission as to inclusion.

Two kinds of silence carry different ethical weight and should not be run together. Technical silence is forced by scale and legend-space limits already licensed by Sections 4.4 through 4.6's trade-off: a road atlas at a scale of one to a million cannot show every footpath, and no particular choice of which minor features to include is thereby political, any more than the choice of projection at that scale is. Motivated silence is different in kind: an omission that tracks something other than legibility, as when informal settlements are left blank on official municipal maps, when a disputed territory is rendered without the boundary line a rival claimant would insist on, or when place names in a minority language are consistently not printed though space for them exists.

This connects to two earlier points in the chapter without collapsing into either. Section 13.5's counter-example established that not every cartographic choice is political; the present distinction supplies the missing other half, since some silences are exactly as innocent as Section 13.5's example and others are exactly what Section 13.4 describes. And motivated silence is the structural inverse of Section 13.6's cartogram, where distortion is declared as the very point of the map: an omission, unlike a cartogram's stated distortion, does not announce itself, and is for that reason harder to audit than a chosen projection or a printed scale bar — the chapter's central claim, that power operates within genuine geometric constraint rather than despite it, extends cleanly to the informational constraint of finite legend space, but only once silence that merely honors that constraint is separated, case by case, from silence that exploits it.

13.9 Latour's Immutable Mobiles: The Map as a Vehicle of Power at a Distance

Chapter 13's discussion so far has concerned what a map states (Sections 13.1–13.7) and what it silently omits (Section 13.8). A further and complementary account, from Bruno Latour's work in science studies, concerns neither content nor omission directly but the map's portability, and what that portability lets an institution do that no single map, considered on its own, would show.

Latour's term for the relevant kind of object is the immutable mobile: a representation able to travel across long distances (mobile) while its internal relations remain fixed in transit (immutable), and able, crucially, to be combined with other such representations once it arrives. Latour's own recurring example is the navigational chart brought back from a voyage to a metropolitan center, where officials who never left port could nonetheless accumulate, overlay, and compare charts from many separate voyages, assembling a picture of distant coastlines more complete than any individual navigator's firsthand knowledge — a form of power Latour calls action at a distance, exercised through paper rather than through presence.

This adds a mechanism to Section 13.4's claim that power exploits the residual freedom Chapter 6 leaves open, rather than duplicating it. Assertion (what a map states) and silence (Section 13.8, what it omits) both operate at the level of a single map's content; the immutable mobile account operates one level up, at the level of format — a stable, combinable representational convention that lets many such maps be centrally collected and set against one another. The same forced geometric trade-off Chapter 4 proves therefore becomes, on this reading, not only a site where a single choice can be politically loaded, but the enabling condition for a center to accumulate comparative knowledge no periphery possesses on its own, moving the world, in Latour's phrase, without moving it.

13.10 Counter-Mapping: The Residual Freedom Turned Against Power

Sections 13.1 through 13.9 have traced how power occupies the residual freedom Chapter 6 leaves open within a genuine geometric constraint — in what a map states, in what it silently omits, and in the format that lets it be accumulated and compared at a distance. The chapter would be one-sided if it closed there, because the same residual freedom Section 13.4 identified as power's point of purchase is, by its own logic, available to anyone who can produce a map, not only to the institutions Sections 13.2 and 13.9 have examined. The scholarly literature has a name for the deliberate use of cartographic convention by the parties a map has traditionally been used against: counter-mapping, a term introduced by Nancy Peluso's 1995 study of community forest mapping in Indonesia and taken up widely since to describe indigenous and local communities producing their own maps, in the same projections, symbol sets, and legal conventions official cartography uses, to document land use, sacred sites, or customary boundaries that official maps had left as Section 13.8 called motivated silence.

The cases collected under this heading share a structure worth stating in exactly Section 13.4's vocabulary. A state cadastral map's omission of an indigenous community's customary territory is, on the diagnosis Section 13.8 already gave, no more forced by Chapter 4's theorem than any other silence; nothing in Tissot's trade-off dictates which boundaries get drawn. Counter-mapping does not dispute this diagnosis so much as act on it directly: a community producing its own GPS-surveyed, professionally drafted map of its own territory is exploiting exactly the same residual freedom — the same underdetermination of center, symbol, and inclusion Section 13.4 located — in the opposite direction, and a number of such maps have gone on to carry legal weight in land-rights litigation and formal title processes precisely because they meet official cartographic convention on its own technical terms rather than rejecting the convention as irredeemably compromised.

The limit on this case is worth stating with Section 13.5's counter-example still in view, so that the chapter's closing note is not read as unearned optimism. Counter-mapping does not escape Chapter 4's theorem any more than the Peters

projection did; a community's map of its own land still trades conformality against area, still centers somewhere and not elsewhere, and still leaves something out that a rival claimant would want shown. What it demonstrates is narrower and, within its scope, secure: that the residual freedom Section 13.4 named is symmetric in a way the chapter's earlier cases, taken alone, might have left the reader assuming it was not — available, in principle, to whoever holds the pen, and not a standing property of institutional power as such. Whether a given community can, in practice, marshal the resources counter-mapping requires — survey equipment, legal counsel, cartographic expertise — is a further and entirely contingent question this volume's mathematics has nothing to say about, and closing the chapter on this note is meant to mark that boundary rather than to blur it.

13.11 A Case in Point: Google Maps and the Border That Changes With the Viewer

Section 13.4 distinguished power's exploitation of the residual freedom Chapter 6 leaves open from any exploitation of the impossibility Chapter 4 proves, and Section 13.8 read a map's silences as a further, subtler exercise of the same residual freedom. A contemporary case combines both moves in a form none of Chapter 13's earlier examples, drawn largely from print-era cartography, had occasion to consider: the same digital map, served by the same company, showing two incompatible political claims as settled fact to two different audiences at once. Google Maps, according to reporting by the Washington Post and earlier accounts from NPR and other outlets, renders several contested borders differently depending on the country from which the map is viewed: Kashmir appears as undisputed Indian territory to users browsing from India and as a disputed region marked by a dashed line to users elsewhere, and the Crimean peninsula appeared, following its 2014 annexation, as part of Russia to users on Google's Russian domain while continuing to appear as Ukrainian territory, separated by a dotted line, to users on Google's international domain. A company spokesperson described the underlying policy as a consistent global

standard of showing disputed regions with a neutral dashed border, combined with compliance to local law on domains explicitly localized to a given country's legislation.

What this adds to Section 13.4's account is not a new instance of power exploiting cartographic discretion, critical cartography's central insight, already well anticipated by Harley's reading of the map as a political act, but a structural variant that insight was not built to describe: a single mapmaker, professing neutrality, resolving the same underlying dispute in opposite directions simultaneously, by segmenting its audience rather than by picking a side. The Peters controversy, Section 13.7 showed, involved a critique that misapplied a real theorem to a single map aimed at a single global audience; the Google Maps case involves no theorem at all, since Tissot's obstruction is silent on how a border, itself not a projection artifact but an attribute layered on top of one, ought to be drawn, and no misapplied proof, only the discretionary residual freedom Section 13.4 identified, now exercised at a scale and with a precision, millions of edits, a great many localized domains, that no print cartographer ever had available. Section 13.10's counter-mapping turned the same residual freedom against power from below; the Google Maps case shows power turning it against itself, maintaining two contradictory maps under one brand rather than committing to either, a possibility this chapter's earlier, print-era cases had no occasion to raise because no print map could ever have been two maps for two readers without either reader knowing it.

Chapter 14 — Conclusion: What Cartography Teaches Every Representational System

14.1 Restating the argument in one page

The argument this volume has tried to keep honest across eleven chapters can be restated briefly, at the cost of the qualifications that gave it its honesty in the first place — qualifications the chapters themselves should be trusted for, not this

summary. A curved surface cannot be represented on a flat one without distortion (Chapters 1 and 4); an everywhere-consistent field of orientation toward a fixed point cannot be defined on a sphere (Chapter 5); the choice among the ways of paying these unavoidable costs is genuinely constrained without being forced, in the specific sense Poincare's conventionalism makes precise (Chapter 6); the first-person structure of needing to face a fixed elsewhere inherits the topological impossibility as a lived, not merely geometric, difficulty (Chapter 7), whose exactness turns out, on inspection of an actual historical practice, to be something a community can require to exist without requiring every member to possess (Chapter 8); even the reference frame within which all of this is stated cannot be assumed to hold still once enough time passes (Chapter 9); and none of this residue of unavoidable cost and constrained choice is itself a flaw to be engineered away, though it is also not, on its own, an alibi for any particular choice made within it (Chapters 10 and 13); the search for a second theorem of Tissot's kind, governing scientific models in general, came back empty, and was reported as empty rather than quietly dropped, even as three independent philosophical accounts of scientific representation converged on the same three-part shape cartography exhibits with proof, and a further, distinct trade-off between descriptive accuracy and explanatory understanding turned out to rest on facts about finite reasoners rather than on the sphere's curvature (Chapters 11 and 12). Fourteen chapters, four theorems, one honestly reported non-result, and one recurring warning against collapsing 'this had to distort somehow' into either 'therefore this distortion was fine' or 'therefore this distortion was chosen for no reason'.

14.2 What travels beyond cartography, and what does not

The Preliminary Note promised to keep two registers apart: results used as rigorous constraints, cited to the chapter of the source volume that proves them, and results used only as image or analogy, marked as such. Closing the volume, it is worth stating plainly which parts of the argument belong to which register, because the two do not travel the same distance. What travels as rigorous

constraint is exactly what Part II and Chapter 9 proved: cartography, specifically, cannot avoid the distortions those chapters name, and no future technique will remove them, because they are theorems, not limitations of technique. What travels only as image or analogy — clearly and repeatedly flagged as such through Parts III and IV — is the wider claim that other representational systems face a structurally similar situation. No general theorem in this volume proves that a legal code, a scientific model, or a person's self-understanding must distort in some analogous, provable sense; the cartographic case is offered as a domain where the structure can be seen with unusual clarity because it can be proven, not as a domain from which the same proof is exported elsewhere. A reader who takes this volume to have proven something about representation in general has taken from it more than it argued, and more than the Preliminary Note said it would try to establish.

14.3 What is left, once the two registers are kept apart

What is left, kept honestly within its proven bounds, is still not nothing. It is the demonstrated existence of at least one domain — a domain of enormous practical and historical consequence, one that people have depended on for navigation, worship, war, and the ordinary business of knowing where they are — in which the wish for a representation with no cost, no residual choice, and no need for a reference frame that holds still is not a wish awaiting better technique but a wish for something that provably does not exist. That existence proof does not by itself settle any question about legal codes, scientific models, or self-understanding. But it removes, for anyone tempted by it, one particular excuse: the thought that unavoidable distortion, contested choice, and an unstable ground are marks of a domain not yet perfected, rather than a possible, unproven, but no longer easily dismissed feature of representation as such. Cartography cannot tell the rest of representation what it is. It can tell it, credibly, what it might be — and where, exactly, to look for a proof.

14.4 A Closing Return to the Preliminary Note's Danger

The Preliminary Note opened this volume with a stated danger: that technical results, once removed from their proofs, are easily turned into loose metaphor, and that a theorem about vector fields on a sphere is not, by itself, a theorem about the limits of any 'global' human orientation, ethical or epistemic. Fourteen chapters later, it is worth asking, as honestly as the volume has tried to ask everything else, whether that danger was in fact avoided or merely stated and then quietly indulged in the writing that followed. The test proposed here is the one Section 14.2 already applied in the abstract: at every point where the argument moved from a companion volume's proof to a claim about facing, dwelling, jurisprudence, power, or representation in general, was the move marked, or did it simply happen. The chapters have tried, subsection by subsection, to mark it: Chapter 7 explicit that the qibla is this volume's chosen instance of a structure the geometry alone does not restrict to religious life; Chapter 9 explicit that the datum's epoch problem is offered as a second, independently interesting case rather than a restatement of the first; Chapter 10 explicit, twice over, about what a proven impossibility does and does not license; Chapter 13 willing, in its own closing section, to find a case where the suspicion its framework invites turns out not to apply.

Whether the marking was sufficient is a judgment this volume cannot make on its own behalf, and does not try to; that judgment belongs to the reader who has followed the argument this far, with the same discipline the volume has asked of itself applied now to the volume itself. What can be said, in closing, is only what the Preliminary Note already promised and what these fourteen chapters have tried, however imperfectly, to deliver: not a general theory of representation dressed in cartography's more presentable, provable clothing, but an honest account of what a small number of genuine theorems, taken exactly as far as they will go and no further, turn out to say about the oldest and most practical of all representational problems, how to put a world too large to hold onto a page small enough to carry, and what is unavoidably lost, and unavoidably chosen, in doing so.

14.5 A Question Left Open, Deliberately: Is There a Meta-Theorem?

The fourteen chapters this volume has built could themselves be asked a question in their own spirit, and it is worth posing before the volume closes rather than leaving it for a reader to notice unasked: is there a meta-theorem here, some further, higher-order impossibility result governing when a domain of representation will or will not admit a Tissot-style theorem of its own? Chapters 11 and 12 searched, honestly and at length, for a second first-order theorem inside the philosophy of scientific models and found none; the question this closing section raises is one level up from that search, not 'does this particular domain have its own impossibility theorem' but 'is there a general theory, provable in its own right, of which domains must have one.' It would be a satisfying way to end this volume if such a meta-theorem existed and could be stated here. It cannot, and the honest thing to do, in keeping with every earlier chapter's practice of reporting a search that came back empty rather than quietly dropping it, is to say so plainly rather than to gesture at a unity the argument has not earned.

What can be said, short of the meta-theorem, is only a pattern noticed across the search rather than a result proved by it: every case in this volume that came closest to a genuine second theorem — the bias-variance decomposition of Section 10.5, the no free lunch theorem of Section 10.7, Batterman's asymptotic explanations in Section 12.2 — turned out, on inspection, to be a statement about a fixed, well-specified mathematical structure, a decomposition of expected error, an average over a stated space of problems, an asymptotic limit, rather than a statement about representation or explanation in some unrestricted, domain-general sense; and every case that stayed merely suggestive, the semantic view's retreat to similarity, structural realism's historical generalization, Simon's bounded rationality, turned out to rest on a judgment, a historical record, or a fact about finite reasoners that no comparable fixed mathematical structure yet specifies exactly. If a meta-theorem is ever found, this pattern is a plausible place to have looked for it: not in representation as such, which is too broad a category for any single proof to range over, but in whether a given

representational practice happens to be reducible to a specific, fully stated mathematical object, the way cartography reduces to projections of a sphere with known curvature. That reduction was cartography's advantage from the volume's first page, and it may, on the evidence this closing section can offer and no more, be the actual precondition any future domain would need to meet before its own Tissot could be found. This is offered, deliberately, as a question for a further volume rather than an answer this one has earned, and it is the last place, before the Conclusion's final paragraph, where the discipline the Preliminary Note asked for still has work left to do.

14.6 What This Volume Deliberately Left Outside Its Scope: Digital, Algorithmic, and Interactive Cartography

Every named projection this volume has used as a case study, Mercator in Section 4.4, the Winkel tripel in Section 4.10, Peirce's quincuncial construction in Appendix B, belongs to a print-era cartography in which a projection is chosen once, by a named cartographer or institution, and fixed for the life of the map. That choice of case studies was not neutral, and the Preliminary Note's discipline of marking what is claimed and what is not requires naming the gap explicitly here rather than letting it pass unmarked into the Conclusion's closing pages. The overwhelming majority of maps now viewed by human beings are not of this kind: they are web maps, served as pre-rendered tiles or generated on demand, almost all of them built on some variant of the Web Mercator projection standardized as EPSG:3857, a single conformal choice adopted not for the reasons Section 4.7 catalogued for print compromise projections, cartographic legibility, areal fairness, aesthetic convention, but very largely because a square, conformal tile can be cut, cached, and reassembled at any zoom level without the seams a non-conformal projection would introduce at tile boundaries, a computational rather than cartographic criterion this volume's case studies never had occasion to consider.

Tissot's theorem, Section 4.1 is careful to note, is a statement about any smooth map from the sphere to the plane, and it does not care whether that map is drawn once by hand or generated ten thousand times a second by a tile server; Web Mercator inherits exactly the distortion Section 4.4 traces in its print ancestor, growing without bound toward the poles, and every zoom level of an interactive map is still, at the level of the individual tile, an instance of the same impossibility Chapter 4 proves. What genuinely lies outside this volume's scope is not the geometry, which travels unchanged, but the philosophical questions a rendering pipeline raises that a fixed print map never could: a projection chosen by an engineering default rather than a cartographer's stated purpose, in a sense Section 6.3's account of purpose-disciplined choice did not anticipate; a map that can, in principle, be re-projected differently for every viewer according to

criteria, a recommender system's inferred intent, a device's screen geometry, that have no analogue among the navigational, thematic, or aesthetic purposes Chapter 4's case studies actually exhibit; and an authorship so distributed, across a projection's original mathematics, a standards body's later adoption, and a rendering engine's runtime choices, that Chapter 13's account of the map as a political act, built throughout around identifiable cartographers and institutions, would need substantial reworking before it could be applied. These are not questions this volume is positioned to answer, and the honest response, consistent with Section 14.2's discipline about what travels and what does not, is to name them as a further volume's proper subject rather than to gesture at answers the fourteen chapters here have not earned.

14.7 A Final Methodological Note: Reporting What Was Not Found

Section 14.5 declined to state a meta-theorem this volume cannot prove, and Chapters 11 and 12 reported at length a search for a second, Tissot-style theorem governing scientific models in general that came back empty rather than quietly setting the search aside. It is worth closing on the method these decisions share, rather than only on their content, because the method is not specific to cartography and is, in a modest sense, this volume's own answer to a question it never posed directly: what does intellectual honesty actually require of an argument built partly on proof and partly on analogy. The answer this volume has tried to enact, chapter by chapter, is that a negative result, honestly reported, is not a failure of the inquiry but a piece of its data, exactly as informative, properly weighted, as a positive one; a reader who finishes Chapters 11 and 12 knowing that no second Tissot theorem was found, and knowing roughly where the search looked and why it came back empty, has learned something a chapter that quietly changed the subject would have withheld.

This is a more general methodological commitment than any single chapter states outright, and naming it here, in the volume's last substantive section before its closing paragraph, is itself an instance of the discipline it describes: results used as rigorous constraint, kept separate from image and analogy, per the

Preliminary Note's opening promise; searches that fail, reported as failures rather than absorbed silently into a rhetoric of success; and a distinction, insisted on from Chapter 1 through the meta-theorem question of Section 14.5, between what a proof has actually established and what a resemblance merely suggests. None of this is original to cartography or to this volume; it is closer to a minimal standard any argument trading on both mathematics and metaphor ought to meet, and stating it explicitly here is offered less as a discovery than as an account of the housekeeping this volume has tried, with whatever success the reader is now positioned to judge, to keep throughout.

14.8 Goodman's Languages of Art, and the Boundary This Volume Has Tried to Hold

The chapters gathered in Part IV have each measured how far a cartographic result travels into a neighboring domain, and each has found a boundary: the bias-variance trade-off (Section 10.5) shares Chapter 4's shape but not its closed-form proof; partial structures (Section 11.2) share Chapter 3's retreat from isomorphism but not its exact classification; Simon's bounded rationality (Section 12.6) shares Chapter 4's trade-off structure but trades off different, non-geometric quantities; critical cartography (Chapter 13) correctly identifies where power operates but, unchecked, mistakes the residual freedom within a constraint for the constraint's absence. One further figure belongs in this company because he represents, more directly than any considered so far, the position this volume has withheld from throughout: Nelson Goodman's Languages of Art argues that pictorial realism is itself a symbol system among others, that a picture strikes a given culture as 'realistic' only to the degree that culture has been habituated to reading its particular representational conventions, and that there is accordingly no privileged, convention-free sense in which one picture resembles its object more than another does — realism, on Goodman's account, is a matter of standard, familiar notation, not of a mind-independent relation of resemblance a picture either achieves or fails to achieve.

Applied to cartography without qualification, Goodman's thesis would dissolve exactly the distinction this volume has spent fourteen chapters insisting on. If which map counts as 'accurate' is wholly a matter of habituated notation, then Tissot's theorem, the hairy ball obstruction, and the datum's dated coordinate are not constraints a map fails or meets; they become, on a fully Goodmanian reading, further conventions among the notational apparatus a culture happens to have settled on, no more binding than the choice of north-up orientation Section 6.7 was careful to classify separately. This is precisely the inflationary temptation Section 1.2 named at the outset, arrived at now by the most sophisticated route available in the philosophy of representation rather than by loose analogy, and it is worth meeting on its own strongest ground rather than dismissing by association with cruder relativisms. What answers it is not a counter-argument internal to semiotics, which this volume has not attempted to adjudicate, but the same fact Section 3.5 already used against the semantic view's most deflationary versions: Tissot's theorem is provable from the differential geometry of the sphere and the plane alone, and holds for every possible projection whether or not any culture has ever looked at a map, exactly as it would hold if every cartographic convention on Earth were replaced tomorrow by an entirely different notational habit. A culture's notational habits can determine which map is read as 'natural' or 'realistic' in Goodman's sense, precisely the phenomenon Section 6.7's Lewis-convention discussion already located and bounded; no culture's notational habits determine whether a given projection is conformal, and no change in habituation could make a projection both conformal and equal-area at once.

This is the boundary the whole volume has tried to hold, stated a final time in its sharpest form: representation is genuinely, and Goodman is largely right that it is, saturated with convention at the level of what strikes a viewer as realistic, familiar, or well-drawn. It is not, and cartography is the discipline able to prove that it is not, thereby saturated with convention at the level of what a given structure-preserving map can and cannot simultaneously achieve. Holding both halves of that sentence together, refusing the deflationary collapse of the second

into the first as firmly as Part I refused the inflationary expansion of cartography's proofs into an unrestricted metaphor, is the single discipline this volume has asked of its reader from the Preliminary Note onward — and the discipline in which, if the fourteen chapters preceding this one have done their work, its whole argument now rests.

14.9 After Paper: Does the Digital Map Change the Philosophical Picture?

The chapters above, including this one's own argument in Sections 14.1 through 14.6, were developed largely against the paper map's fixed, single, one-scale-at-a-time projection. Before closing, it is worth asking directly whether the continuously zooming, server-rendered, often personalized digital map that constitutes most ordinary map use today unsettles any of this volume's conclusions, or only their most familiar illustration.

The trade-off itself does not move. Chapter 4's theorem constrains any single flat rendering at any given moment, and a slippy digital map does not escape this so much as re-solve the projection choice at each zoom level, typically defaulting to a local variant of the Web Mercator projection for reasons Chapter 6's conventionalism would recognize at once — tiling convenience and locally acceptable angular fidelity, not any claim to global equal-area truth. What does change is Section 13.8's silence and Section 3.6's discussion of the map's fictional, DEKI-theoretic status: a digital map's visible content can be selectively rendered per query, per user, or per paying advertiser, so that the single stable object the earlier chapters sometimes assumed gives way to a family of momentary, individually generated renderings, each one separately answerable to Chapter 4's trade-off, none of them any longer 'the' map in the older, single-artifact sense Borges's fable (Section 2.9) still took for granted.

This closes in the register the Preliminary Note set at the outset. Nothing here is offered as a new theorem, and the volume has none to offer on this point; it is the observation that the conditions making cartography philosophically instructive

in the first place — a forced, provable trade-off standing behind an ordinary, practical object — survive the migration from paper to screen fully intact, while the object itself, and with it the politics Chapter 13 examined and the phenomenology Chapter 7 examined, do not.

14.10 A Final Return to Borges: The Fable Cartography's Proofs Make Literal

Section 2.2 opened this volume's genealogy with Borges's one-paragraph fable of an empire whose cartographers, dissatisfied with every map short of perfection, eventually produce one at the scale of the empire itself, a map coextensive with the territory it represents and consequently useless, abandoned to the elements by the following generation. Section 2.9 returned to the fable once more, taking its premise literally enough to ask what a genuine one-to-one correspondence between map and territory would actually require, and found the demand incoherent for reasons closer to Cantor than to Borges's own gentle satire. Twelve chapters later, with Chapter 4's theorem proved, Chapter 5's second impossibility proved, and Chapter 9's datum instability established as a third, it is worth asking a question the volume's opening genealogy could not yet answer: what, exactly, would the empire's cartographers have discovered had they stopped one step short of Borges's absurd conclusion and instead simply kept enlarging an honest, non-literal map, seeking not identity with the territory but merely to drive distortion toward zero by making the map ever larger and the region it covered ever smaller.

The answer, available only now that Part II's proofs are behind the volume rather than ahead of it, is that they would have found Section 4.8's Chebyshev result waiting for them: for a sufficiently small region, a conformal map's distortion can indeed be driven arbitrarily close to zero, and the empire's cartographers, had they wanted merely a good local map rather than a perfect universal one, could have gotten arbitrarily close to what Borges's fable treats as available only at the absurd limit of a map the size of the empire itself. What they could not have gotten, at any finite scale short of that absurd limit, is exact fidelity over an

extended region, because Chapter 4's theorem, not merely Borges's satire, forbids it; the fable's joke, that only a map as large as the territory achieves perfect correspondence, turns out on this volume's evidence to contain a truth its author was not arguing for and very likely did not have in mind, since Gauss's proof predates the 1946 fable by more than a century but was, so far as the volume's research for Chapter 2 could establish, no part of Borges's own frame of reference. Borges reached, by literary instinct, the same limiting case Chapter 4 reaches by proof: that perfect cartographic fidelity is available only in a case degenerate enough to be useless, a one-paragraph fable and a chapter-length theorem converging, from opposite directions and for entirely unrelated reasons, on the same unforgiving point. It is a fitting note on which to let this volume's genealogy close, twelve chapters after Section 2.2 first told the story without yet being able to say why it was true and not merely clever.

14.11 A Last Instruction to the Reader: How to Use What This Volume Has Not Settled

Fourteen chapters have accumulated, alongside their proofs, a substantial ledger of marked uncertainty: a meta-theorem Section 14.5 explicitly declined to state, a scope Section 14.6 named rather than covered, a search Chapters 11 and 12 reported as having come back empty, and a great many places throughout where a formal parallel was offered as suggestive rather than binding. A reader reaching this point with all of that still in view might reasonably ask what to do with an argument built this way, proved in places and merely gestured at in others, rather than uniformly certain from the first page to the last. The answer this volume offers, as its genuinely final instruction, is to keep the ledger rather than resolve it: to use Chapter 4's theorem, Chapter 5's theorem, and Section 9's account of datum instability as load-bearing, exactly as certain as their proofs make them, and to use everything built by analogy around them, Chapters 6 through 13's extensions into convention, dwelling, orientation, ontology, distortion, and power, as illuminating without being thereby proved, a

distinction this volume has tried to mark at every place it mattered rather than only in its final pages.

This is not a counsel of skepticism, and it is not, despite the discipline it asks for, a hedge against being wrong; it is closer to the only honest way to hand a reader an argument that mixes genuine mathematics with genuine metaphor without letting the metaphor borrow the mathematics's certainty. A cartographer's map is not the territory, Chapter 2 established at length, but a good map, honestly labeled as to where its own distortions lie, is still worth using to get somewhere; this volume has tried, chapter by chapter, to be exactly that kind of map of its own subject, distorted in places it has tried to name, and offered, on that condition, as worth the reader's trust regardless.

A Note on the Appendices

The fourteen chapters above are the complete argument this volume set out to make, and the Conclusion's summary in Section 14.1 stands as a summary of exactly those fourteen chapters, no more and no fewer; nothing that follows is required reading for that argument, and nothing that follows revises it. The three appendices that follow instead do what an appendix in this volume's tradition of citing companion volumes has always done: they take a case mentioned in passing in the main text — discrete global grid systems, noted only as the subject of a companion volume in the Preliminary Note; the named projections of Section 4.4 and 4.5; the three retroazimuthal constructions named without individual treatment in Section 5.2 — and give it the chapter-length attention the main argument's pace could not afford without losing its own thread. Each appendix applies the same discipline the main chapters applied: a rigorous constraint is distinguished from an analogy, and no claim is made that outruns what is actually shown.

Appendix A — Discrete Global Grid Systems and the Continuous/Discrete Divide

A.1 A Different Kind of Reduction

The chapters of Part II treated cartography's distortion as a fact about continuous surfaces: a curved manifold, mapped by a smooth function, onto a flat plane, with the distortion at each point given by the local scale factors of Section 4.1. A companion technical volume, From Continuous to Discrete, examines a different family of representations that this volume's main argument has, until now, left entirely to one side: discrete global grid systems, DGGS, which do not map the sphere onto a continuous plane at all, but instead partition it directly, at the outset, into a finite mesh of discrete cells — hexagons, triangles, or quadrilaterals, tiled recursively at successively finer resolutions — and represent every location not by a continuous coordinate pair but by the identifier of the cell containing it. Systems such as H3, S2, and the rHEALPix family, worked with computationally in the companion volume, differ among themselves in cell shape and in the recursive subdivision scheme each adopts, but share the same basic move: geographic information is stored and queried cell by cell, not point by point, and the question a Tissot-style theorem asks of a continuous projection — how much does a small neighborhood's shape or area distort — has no direct analogue for a system that never claims to represent a neighborhood smaller than a single cell in the first place.

A.2 Why This Is Not an Escape from Chapter 4's Theorem

It would be a natural, and mistaken, thought that a discrete grid sidesteps Tissot's theorem by simply declining to make the continuous claim the theorem constrains — no cell boundary claims to preserve angle or area at a point, so no theorem about angle and area at a point can be violated. This is true as far as it goes and considerably less interesting than it first appears, because the discrete representation inherits its own version of the same constraint one level up: a

DGGS built from cells of equal area on the sphere cannot, in general, also make every cell equilateral or angle-regular in its planar unfolding, and a DGGS built for uniform cell shape typically sacrifices exact area equality among cells at different latitudes — the same trade-off Chapter 4 proved for continuous projections, now stated as a property of a finite partition rather than of a differentiable map, and traceable to the identical fact, the sphere's curvature, that forces the continuous case. The companion volume's own computational comparisons among H3's hexagonal cells, S2's quadrilateral cells, and rHEALPix's equal-area cells document exactly this: no scheme achieves both uniform cell shape and uniform cell area at every resolution and every latitude, and the schemes differ mainly in which of the two a given application needs preserved, not in whether the trade-off can be avoided.

A.3 What Discreteness Adds to the Philosophical Picture

What is genuinely new in the discrete case, and not already covered by Chapters 3 through 6, is a further reduction the continuous case does not force at all: quantization, the collapse of every location within a cell to a single represented value, the cell's identifier, regardless of where within the cell the true location falls. This is a homomorphism in Section 3.4's exact sense — many empirical locations mapped to one representational symbol — but of a different character than a projection's distortion, because it is chosen, deliberately, at a stated resolution, rather than forced continuously at every point; a DGGS user who needs finer distinction simply selects a finer resolution level, trading storage and computation for reduced quantization error, a choice with no analogue in Tissot's theorem, where no choice of scale removes the distortion at a point. The discrete case thus layers a second, genuinely optional reduction, quantization, resolution-dependent, on top of a first, forced one, the area/shape trade-off inherited from curvature, and keeping the two apart, one a matter of Chapter 6's residual freedom, the other a matter of Chapter 4's forced cost, is exactly the discipline this volume has asked of every case it has examined, applied here to a representational form its main chapters never directly treated.

A.4 The Twelve Pentagons: A Discrete Echo of Chapter 5's Parity Argument

Section 5.5 asked why the hairy ball obstruction forces exactly two singularities rather than one or twenty, and answered from the Poincaré–Hopf theorem: the indices of a vector field's zeros must sum to the sphere's Euler characteristic, two, and a single index-two zero is topologically fragile, generically splitting into the familiar pair. Discrete global grid systems built by recursively subdividing an icosahedron—H3 among them, as Section A.1 already named—exhibit a close discrete cousin of exactly this forcing, worth stating because it is easy to miss amid Section A.2's already-established point that discreteness does not evade Chapter 4's continuous trade-off. A hexagonal tiling of the plane can be extended indefinitely using only hexagons, each vertex meeting exactly three others at the hexagon's interior angle; the sphere's curvature forbids this. Euler's formula, $V - E + F = 2$, applied to any trivalent polyhedron built predominantly from hexagonal faces, forces exactly twelve faces to be pentagonal rather than hexagonal, regardless of how many hexagons surround them or how finely the whole structure is subdivided—the same twelve-pentagon count familiar from the geodesic domes and Goldberg polyhedra Section 5.1 already mentioned in passing, and from an ordinary soccer ball's stitching pattern.

The parallel to Section 5.5's argument is structural rather than a claim that the two theorems are the same result: Poincaré–Hopf constrains a continuous vector field's zeros by an integer invariant of the sphere; Euler's formula constrains a discrete trivalent tiling's face count by the same invariant, applied combinatorially rather than differentially. Both are, in the precise sense Section 5.8 used for Arrow's theorem, proofs that a global requirement (an everywhere-hexagonal tiling; an everywhere-nonvanishing field) cannot be met without a forced, quantifiable exception, traceable in each case to the sphere's Euler characteristic rather than to any deficiency of construction. H3's twelve unavoidable pentagonal cells are, in this sense, not a defect of that particular indexing scheme to be engineered away by a cleverer subdivision; they are the discrete, combinatorial shadow of the identical topological fact Chapter 5 proves

continuously, appearing here in a representational form Part II's continuous chapters had no occasion to exhibit directly.

What this adds to Appendix A's own argument is a second instance of Section A.2's claim that discreteness does not exempt a representation from curvature's consequences, now at the level of topology rather than metric distortion: a DGGS cannot only fail to achieve uniform cell shape and uniform cell area simultaneously (Section A.2's trade-off); a hexagon-based DGGS cannot even achieve uniform cell shape on its own, at any resolution, because twelve cells at every subdivision level are forced to be a different shape from their neighbors. Practitioners using H3 accommodate this in practice by placing the twelve pentagons at fixed, documented locations chosen to fall over ocean where possible, a design choice inside Chapter 6's residual freedom exactly as Section 4.3 described a cartographer choosing where a projection's distortion falls—the forcing is topological and inescapable, but where the twelve unavoidable exceptions are placed remains exactly as free, and exactly as answerable to a stated purpose, as any other choice this volume has examined.

A.5 Space-Filling Curves: An Optional Layer Above the Forced One

Section A.3 distinguished a forced reduction, the area/shape trade-off inherited from curvature, from an optional one, quantization to a chosen resolution, and located discreteness's genuine philosophical novelty in the second. A further layer, orthogonal to both and worth adding for completeness, concerns not what a DGGS represents but how its cells are ordered once they must be stored, indexed, or transmitted linearly—a practical requirement Section A.1's cell-by-cell query model does not by itself resolve, since a cell identifier alone fixes which region a cell covers but not where that identifier falls in a sorted list relative to its geographic neighbors. Space-filling curves, orderings such as the Hilbert curve or the simpler Morton (Z-order) curve, solve this by mapping multidimensional coordinates to a single linear index in a way that keeps points close in space close, though not perfectly close, in the resulting order, and

systems in the companion volume's comparison, S2 among them, build their indexing directly on such a curve over the sphere's recursive subdivision.

This choice is optional in exactly Section A.3's sense, and worth distinguishing sharply from both of Appendix A's other two reductions so the count of distinct layers stays exact. The area/shape trade-off (Section A.2) is forced by curvature and admits no choice at all; quantization (Section A.3) is a choice of resolution, trading precision against storage, but every resolution still requires some cell to be queried directly; the choice of space-filling curve is narrower still, a choice of enumeration only, changing nothing about which cell a location falls in or how finely the sphere is subdivided, and affecting only how efficiently a range of nearby cells can be retrieved as a contiguous block of the linear index. A Hilbert curve and a Morton curve applied to the identical H3 or S2 subdivision produce the identical partition of the sphere and the identical quantization error at every resolution; they differ only in how well geographic locality survives the flattening into a single sorted list, the Hilbert curve generally preserving it better at the cost of a more complex encoding and decoding computation.

The lesson this closing addition draws out, and the reason it belongs in this appendix rather than as a footnote to Section A.3, is that even after Chapter 4's forced cost and Section A.3's optional quantization have both been paid, a representation still carries further, genuinely free design choices with no analogue anywhere in the continuous case Part II examines—because a continuous projection has no corresponding need to linearize its points into a storage order at all. Discreteness does not merely add one new reduction to cartography's inherited one, as Section A.3 already showed; it opens a further design space, layered on top of both, that is a genuine artifact of representing a curved world inside a machine built to store its data one address at a time.

Appendix B — The Peirce Quincuncial Projection: Conformality Pushed to Its Limit

B.1 Taking the Trade-off at Its Word

Chapter 4 established that a projection may choose conformality at the cost of area, or area at the cost of conformality, or some compromise between the two; the named cases examined there, Mercator, Albers, Robinson, and Fuller's Dymaxion in Section 4.5, occupy different points along that space without any of them pressing all the way to conformality's own limit. Charles Sanders Peirce's quincuncial projection, published in 1879, is worth a chapter-length look precisely because it does press to that limit, in a construction unrelated in its mathematics to any of Chapter 4's cases: built from a stereographic projection of the sphere onto the plane, itself already conformal, mapping angles perfectly but sending area toward infinity near the projected pole, composed with a further conformal map, given by an elliptic function, that folds the infinite stereographic plane into a finite square. The construction preserves conformality at every point of the composition, because a composition of conformal maps is conformal, and the square that results tiles the plane seamlessly when repeated, a property with no analogue in Mercator's chart, which cannot be tiled without a seam.

B.2 The Price Paid, and Where

Total conformality, on Chapter 4's own terms, cannot be free; the theorem does not permit an exception even for so ingenious a construction, and the Peirce projection's price is paid in two places the named cases of Section 4.4 do not exhibit. First, area distortion is not merely present but extreme and non-uniform, growing without bound toward the branch points the elliptic function introduces at prescribed points of the square, a direct, unavoidable consequence of insisting on conformality everywhere, exactly as Section 4.1's statement of the theorem requires. Second, and more distinctive to this construction than to any case Chapter 4 treated, the branch points themselves are not ordinary points of the

map at all: the elliptic function's structure forces a specific angular deficiency at each one, so that the map is conformal in the strict sense at every point except a finite set where even that guarantee fails in a controlled, mathematically characterized way — a residual singularity structurally reminiscent of, though mathematically unrelated to, the two singularities Chapter 5 proved must attend any global field of fixed directions, and worth flagging as a reminder, in the spirit of Section 5.3's caution, that the resemblance is one of shape, a global requirement forcing a finite set of unavoidable exceptions, and not one of cause: Peirce's branch points are forced by the choice of an elliptic function for a conformal folding, not by the hairy ball theorem, and the two impossibilities, though visually and structurally alike, belong to different branches of mathematics and must not be allowed to borrow each other's proof.

B.3 An Aesthetic Criterion for the Residual Freedom

The Peirce projection has had, historically, almost no practical cartographic use of the kind Section 4.3 organized its discussion around — no navigator needed it, no demographer chose it — and its afterlife has been almost entirely in a different register: as an object of mathematical and aesthetic interest, admired for the seamless tiling its conformality makes possible and reproduced, since the twentieth century, in decorative and artistic contexts having nothing to do with navigation or thematic mapping at all. This is worth setting beside Chapter 6's account of the residual freedom the impossibility theorems leave open: Section 6.3 argued that a purpose disciplines the choice among survivors once the theorem's cost is known, and the Peirce case shows a purpose Chapter 6 did not explicitly list among its examples, elegance or formal interest for its own sake, functioning exactly as navigational or thematic purpose functions elsewhere in this volume — a criterion that selects one point in Tissot's constrained space over another, answerable, as every such choice is, to the question of what it is for, even where the answer is not utility in any of the ordinary senses Chapter 4 organized its case studies around.

B.4 The Adams and Guyou Projections: A Small Family, Not a Single Curiosity

Appendix B has so far treated Peirce's quincuncial projection as though it stood alone, an isolated extremum of Chapter 4's trade-off space; it belongs, in fact, to a small family of related conformal square constructions, and naming the family sharpens rather than dilutes Section B.1's point about what it means to press conformality to its limit. Émile Guyou published a conformal projection of the sphere onto a square in 1887, predating Peirce's own publication's wider circulation, oriented so the square's diagonals align with the equator and the central meridian rather than, as in Peirce's construction, with the poles; Oscar Adams, working at the U.S. Coast and Geodetic Survey, published a further variant in 1925, the Adams hemisphere-in-a-square projection, adapted to map a single hemisphere rather than the whole sphere into one square, each of the three built from the identical mathematical apparatus Section B.1 described, a stereographic projection composed with a conformal elliptic-function folding, differing among themselves only in the placement of the singular points the folding introduces.

The family relation matters for the same reason Section 4.5's Dymaxion case mattered to Chapter 4: it distinguishes a genuine degree of freedom, still governed by the same forced cost, from an isolated trick exempt from it. Guyou's and Adams's projections pay Section B.2's price in exactly the form Section B.2 describes—extreme, non-uniform area distortion and a forced finite set of branch-point singularities—differing from Peirce's own construction only in where, relative to the geographic graticule, that price is paid, a choice inside Chapter 6's residual freedom in precisely the sense Section B.3 already argued for Peirce's case alone. What the family adds is evidence, of the same modest kind Section 5.7 drew from nematic liquid crystals, that conformality's extreme point is not a single mathematical curiosity but a genuine region of the trade-off space, reachable by more than one construction and differing among its instances only in the free choice of where the singularities fall, not in whether they are forced to exist at all.

B.5 Escher, Coxeter, and a Conformal Aesthetic Beyond the Sphere

Section B.3 located the Peirce projection's afterlife almost entirely in an aesthetic register, valued for the seamless tiling its conformality makes possible rather than for any navigational or thematic use. A related but geometrically distinct case, worth adding as the appendix's closing extension in the spirit Section 5.7 modeled for Chapter 5, comes from a surface with the opposite curvature to the sphere's: M.C. Escher's Circle Limit series, produced from the mid-1950s after a documented correspondence with the geometer H.S.M. Coxeter, renders repeating tilings of the hyperbolic plane, using the Poincaré disk model—itself a conformal map, exactly as Section B.1's stereographic projection is, preserving angle while letting scale shrink without bound toward the disk's boundary, so that infinitely many congruent tiles, identical in the hyperbolic metric, appear to a Euclidean eye to shrink toward the rim.

The disanalogy is worth stating with the same care Section B.2 gave the hairy-ball resemblance, because it is easy to let two conformal, tile-generating constructions blur into one loosely shared idea. Peirce's square tiles a surface of constant positive curvature, the sphere, and the branch-point singularities Section B.2 traces are a direct consequence of that positive curvature meeting a Euclidean square's corners; Escher and Coxeter's disk tiles a surface of constant negative curvature, and its infinite, boundary-crowding repetition is a signature of hyperbolic rather than spherical geometry, with no branch points of Peirce's kind anywhere in its construction, precisely because the hyperbolic plane, unlike the sphere, admits an isometric-enough conformal model in the first place for this specific decorative purpose. What the two cases share, and the only claim this section makes on their behalf together, is Section B.3's already-established point generalized one step further: that a conformal map's value need not be cartographic at all, and that mathematicians working on surfaces of entirely opposite curvature, a nineteenth-century logician's sphere and a twentieth-century artist's hyperbolic disk, converged independently on the same non-utilitarian criterion, elegance and tileability, for judging what a conformal construction is for.

B.6 A Twenty-First-Century Coda: Conformal Squares in Digital Rendering

Sections B.1 through B.5 have traced Peirce's construction and its nineteenth- and twentieth-century relatives, Guyou's and Adams's variant foldings, Escher and Coxeter's hyperbolic analogue, as a family united by a non-utilitarian criterion, elegance and tileability, standing outside the navigational and thematic purposes Chapter 4 organized its main case studies around. That family has acquired, in the past two decades, a coda worth recording briefly rather than leaving the appendix to close in the nineteenth century its central case belongs to. The elliptic-function computations that made Peirce's, Guyou's, and Adams's foldings a matter of considerable hand calculation in their own time, John P. Snyder's standard reference on map projection formulas devotes several pages to approximating the relevant elliptic integrals precisely because closed-form evaluation was impractical without a computer, are now a matter of a few dozen lines in any general-purpose projection library, most widely PROJ, the open-source engine underlying the majority of contemporary geographic information systems. What was, for Peirce and Guyou, a demonstration of what conformality could achieve at the cost of considerable analytical labor is, for a present-day user of such a library, a keyword argument.

This computational trivialization is worth setting beside Section B.3's aesthetic criterion rather than treated as a separate development, because it has not displaced that criterion but extended its reach. The same low cost of computing a conformal square folding has made the Peirce family newly available to procedural and generative digital artists, who use it, as Escher used the Poincaré disk by hand in Section B.5's account, to produce seamlessly tiling imagery for its own visual interest, now at a scale and iteration speed no nineteenth-century construction by hand permitted; websites and code libraries built specifically to generate quincuncial map art, unconnected to any cartographic institution or navigational purpose, are direct descendants of the criterion Section B.3 identified in Peirce's own nineteenth-century afterlife. The appendix's closing observation is accordingly a modest one: the trade-off Chapter 4 proves is

indifferent to the computational cost of standing at any particular point within it, but the ease or difficulty of reaching a given point has its own history, and the Peirce family's history runs from a nineteenth-century logician's labor-intensive elegance to a twenty-first-century default library call without the underlying mathematics, or the underlying price Section B.2 traces, changing in any respect at all.

Appendix C — The Three Retroazimuthal Constructions as Aesthetic Objects

C.1 Three Accommodations of One Obstruction

Section 5.2 named, without dwelling on them individually, the three historical cartographic solutions to the retroazimuthal problem: the Littrow projection, the Craig retroazimuthal projection, and the Hammer retroazimuthal projection, each accommodating the hairy ball obstruction Chapter 5 proves cannot be defeated on its own terms. It is worth pausing on what, exactly, each gives up, because the three do not make the same concession, and the differences among them are a second, smaller instance of Chapter 4's trade-off logic operating inside a situation Chapter 5 already established admits no trade-off at the level of the obstruction itself. The Littrow projection preserves the correct azimuth to a single chosen reference point along an entire meridian through that point, at the cost of conformality away from that meridian and of any claim to accuracy for bearings to points other than the one the meridian is built around. The Craig retroazimuthal projection, designed originally to show the qibla direction from a single city, gives correct azimuth to its reference point from every location on the map, a stronger property than Littrow's, purchased at the cost of extreme, rapidly worsening area and shape distortion away from the reference point and its immediate surroundings. The Hammer retroazimuthal projection modifies the underlying construction again, trading a different combination of the same currencies, correct bearing to the reference point against area fidelity, along a different curve of compromise than Craig's.

C.2 Elegance Internal to a Forced Concession

None of the three constructions is more correct than the others in any sense Chapter 5's theorem settles; the theorem fixes only that some concession is unavoidable, not which one, and the three historical solutions are, in this precise sense, three different answers to a question Chapter 6's conventionalism already prepared this volume to expect would need an answer beyond the geometry alone. What is worth adding here, as this volume's closing extension of that discussion, is that the choice among the three has been governed, historically and among specialists in the projection's own small literature, at least as much by a sense of mathematical elegance, the simplicity of the construction, the tractability of the resulting formula, the visual regularity of the distortion pattern, as by any difference in practical utility among the reference points each construction happens to have originally been built around; a reference point could, in principle, be swapped in any of the three constructions without changing which cartographer might prefer it, and preference has tended to track the construction's own formal character rather than the practical stakes of any one reference point over another.

C.3 What This Adds, and Where the Discipline Still Applies

Naming this preference aesthetic rather than practical is not a criticism, and it is not, either, evidence of anything Chapter 5's theorem itself did not already anticipate: Section 5.4 already framed Part III's task as asking what it is to live inside a constraint the geometry proves cannot be relaxed, and a cartographer's preference for one accommodation's elegance over another's is one more instance of exactly that living, no more exempt from Chapter 6's requirement that a constrained choice still owe an account of itself than the political choices Chapter 13 examined at length. The discipline this volume has tried to hold throughout applies here in its smallest and least consequential form: an aesthetic preference among three mathematically equal accommodations of a proven impossibility is a real and legitimate criterion for choice, exactly as Appendix B argued for the Peirce projection's afterlife, but it remains what Chapter 6 always said the

residual freedom's content must be, a further, non-geometric reason, owed to the reader once the theorem's own silence on the matter is acknowledged, and not a fact the geometry itself could ever have supplied.

C.4 The Reference Point as a Free Parameter, Not a Fixed Feature

Sections C.1 through C.3 have discussed Littrow, Craig, and Hammer's constructions largely through their best-known historical association, Craig's in particular with a single reference point, Mecca, that motivated its original design. It is worth making explicit a distinction the discussion so far has left implicit, because Chapter 8's treatment of the qibla specifically could otherwise leave the impression that Craig's construction is somehow built for that one city rather than for the general retroazimuthal problem Chapter 5 states impersonally. Each of the three constructions is, mathematically, a function of an arbitrarily chosen reference point C, not a fixed feature tied to any one city; recentering Craig's projection on London, or Tokyo, or the antipode of Mecca itself, changes the numerical coefficients of the construction but not its structure, its singularities, or the shape of the concession Section C.1 already described it as making.

This sharpens Section C.2's claim that the historical choice among the three constructions has tracked elegance rather than the practical stakes of any one reference point, by making clear why that could be so at all: since every one of the three accommodates the identical Chapter 5 obstruction for any C whatsoever, a cartographer choosing among Littrow, Craig, and Hammer for a newly commissioned map, oriented toward some other city entirely, faces exactly the same three-way aesthetic choice this appendix has already described, undiminished by the fact that each construction happens to be best remembered for the one historical instance that first motivated it. The reference point is a parameter the mathematics leaves open in exactly Chapter 6's sense; the fame of one particular value of that parameter, Mecca's, is a fact about the history of the qibla problem specifically, examined at length in Chapter 8, and not a fact about what the three constructions themselves require.

C.5 A Computational Coda: When the Aesthetic Question Stops Being Asked

Sections C.1 through C.4 have treated the choice among Littrow, Craig, and Hammer as a live aesthetic question, precisely because each is a static, printed construction, a single fixed image that must accommodate Chapter 5's obstruction once and for all across its entire printed extent. Section 7.7 already noted, in passing, that a modern smartphone compass app restores something close to a genuine global bearing field computationally, for a single reference point, at the cost of the device the field is displayed on rather than at the cost of a printed projection's shape; it is worth closing this appendix by stating plainly what that observation implies for the aesthetic question Sections C.1 through C.4 have just spent themselves on. A device that recomputes the bearing to Mecca, or to any other C, freshly from the user's current position each time the question is asked never needs to choose among Littrow's, Craig's, or Hammer's concessions at all, because it never produces a static map of the whole sphere's bearings in the first place; it answers Section 7.2's question, 'which way must I turn from here', pointwise, for the single P the user actually occupies, and simply recomputes when P changes, leaving the two singularities Section 5.5 identified exactly where the theorem places them, at C and at its antipode, but rendering the question of how to accommodate them across an entire printed surface moot for a user who only ever needs the answer at one point at a time.

This does not retire Chapter 5's theorem, nor Chapter 6's conventionalism, nor this appendix's own aesthetic argument; a device answering pointwise still fails, undefined, exactly at C and at C's antipode, and a cartographer asked to produce a printed wall map for a mosque, rather than a phone screen for a single traveler, still faces precisely the three-way choice Sections C.1 through C.4 describe. What the computational case adds, and the reason it belongs as this appendix's final word, is a sharper sense of which of this volume's cases are artifacts of a specific representational technology, the single static printed sheet Part II's entire argument has largely presupposed, and which are forced by the geometry itself, independent of whatever technology renders it: Chapter 5's two singularities

belong to the second kind absolutely, surviving the move from paper to computation unchanged, while the aesthetic question this appendix has posed throughout, which static concession looks best, turns out to belong partly to the first, a question a sufficiently different representational technology can render differently, or, as here, simply stop needing to ask.

Appendix D — Toponymy and the Entropy of Naming: The Symbolic Layer Cartography Presupposes

D.1 A Third Layer, Neither Continuous Nor Discrete

Chapter 4 treated distortion as a fact about continuous surfaces; Appendix A treated a second reduction, quantization, forced by the choice to partition the sphere into discrete cells rather than map it continuously. A third layer of cartographic representation has been present in every map this volume has discussed and examined in none of them: the names written on it. A place-name is neither a continuous distortion of the kind Chapter 4 proves inevitable nor a discrete quantization of the kind Appendix A describes; it is, in the vocabulary Section 3.8 borrowed from Peirce, a symbol in the strict sense, a sign connected to its object by convention rather than by resemblance or causal trace, and the companion technical volumes in this author's toponomastic series, the Entropy volumes and the Networks volumes built on them, treat the resulting system of names, across substrate zones defined by shared linguistic origin rather than by political boundary, with the same quantitative rigor Chapters 4 and 5 bring to angle and area.

The method, briefly: a large corpus of attested toponyms, drawn from gazetteer data for a given region, is treated as a sample from a distribution over bigrams, the two-letter sequences a naming system favors or avoids, and the Shannon entropy of that distribution is a genuine, computable quantity, a measure in bits of how much a name's letter sequence, on average, actually distinguishes it from its neighbors within the local naming tradition. Comparing two adjacent

substrate zones' bigram distributions with the Jensen-Shannon divergence, a symmetrized, bounded measure related to but better-behaved than the Kullback-Leibler divergence it is built from, yields a number for how different two naming traditions are at the level of their letter statistics, independent of any single etymology; treating the square root of that divergence as a Fisher-Rao-related distance allows the substrate zones to be arranged, by standard clustering methods, into a dendrogram, and the resulting tree structure to be compared, quantitatively, against the actual geographic contiguity of the same zones on the ground.

D.2 An Entropy Cost With No Tissot Analogue, and One That Does

Two distinct findings from the companion series are worth separating with the same care Chapter 4 and Appendix A separated a forced cost from an optional one, because they do not have the same epistemic status and the difference matters. Bigram entropy itself, the information carried per name within a single substrate zone, is not a distortion in Section 4.1's sense at all; it measures how much a naming system distinguishes, not how well a naming system corresponds to anything geometric, and no theorem of this volume's cartographic chapters bears on it one way or the other. Where the toponomastic material does meet this volume's argument is at the second step: when the Networks volumes compare the dendrogram built from inter-zone divergence against the actual adjacency graph of the same zones on the map, a Declared Geographic Contiguity Graph checked against a Louvain community partition of the divergence data, the two structures agree only partially, with an adjusted Rand index and normalized mutual information both well short of the value that would indicate perfect correspondence.

This partial correspondence is the toponomastic layer's nearest analogue to the impossibility results Chapters 4 and 5 prove, and the disanalogy is exactly as important as the analogy. Tissot's theorem and the hairy ball theorem are proven: no projection, however constructed, and no vector field, however clever, can escape them, for reasons available in closed form from the differential geometry

alone. The imperfect match between linguistic-entropic distance and geographic contiguity is not proven in that sense at all; it is an empirical, statistical finding about the specific corpora the companion volumes have actually assembled, and a different or larger gazetteer, a different bigram window, or a different clustering rule could in principle narrow or widen the gap the Networks volumes report, in a way no re-measurement could ever narrow Tissot's trade-off. Whether some version of the mismatch is forced, the way curvature forces cartographic distortion, or is merely a contingent feature of how human naming traditions happen to have diffused across contact zones, is a question this volume's technical companions leave, honestly, open.

D.3 Why This Belongs in an Appendix, Not a Chapter

That open question is the reason this material appears here, as an appendix, rather than as a chapter within the volume's main argument, and the reason is worth stating rather than leaving implicit, since the volume has asked its reader throughout to notice exactly this kind of distinction. Chapters 4 through 6 and Chapters 7 through 9 rest on results that are proven: a projection cannot be both conformal and equal-area, a tangent field on the sphere cannot be everywhere nonvanishing, a geodetic datum cannot be held fixed indefinitely without a published epoch and velocity model, and every chapter built on these results inherits their necessity. The toponomastic findings summarized here are of a different kind, rigorous within their own statistical framework, honestly and extensively cross-validated across more than thirty substrate zones in the companion volumes, but empirical rather than necessary, descriptive of a body of gazetteer data rather than derivable from geometry alone. A chapter of this volume's main sequence claims, implicitly, the same standing as its neighbors; an appendix does not, and placing the toponomastic material here, alongside the discrete grids of Appendix A and the aesthetic afterlives of Appendices B and C, is this volume's way of giving the naming layer of a map its due, as a genuine third mode of representation worth a rigorous look, without borrowing for it a certainty its own method does not, and does not claim to, supply.

D.4 Zipf's Law: A Further Regularity the Bigram Model Does Not Capture

Section D.1 described the toponomastic method as a bigram-level analysis: attested place-names treated as a sample from a distribution over two-letter sequences, compared across substrate zones by the Jensen–Shannon divergence Section D.1 already introduced. A second, independent regularity is well documented in the general linguistic literature on naming frequency, worth noting here precisely because it operates at a different level than anything Section D.1's method touches: George Kingsley Zipf's observation, first made for word frequency in running text and repeatedly confirmed since across many kinds of naming systems, that when items in a corpus are ranked by frequency of occurrence, frequency tends to fall off as an inverse power of rank, a small number of names accounting for a disproportionate share of all occurrences while a long tail of names each occur only once or a handful of times in the gazetteer.

This is worth flagging as a distinct finding rather than folded into Section D.1's entropy measure, because the two describe genuinely different properties of the same corpus and neither determines the other. Bigram entropy, Section D.2 already argued, measures how much a name's internal letter sequence distinguishes it from its neighbors within a naming tradition; a Zipfian rank-frequency profile measures something orthogonal entirely, how unevenly the names in a gazetteer are distributed across the landscape's actual instances, generic terms such as a local word for 'stream' or 'hill' recurring across hundreds of features while distinctive commemorative or descriptive names occur once. A substrate zone could in principle exhibit high bigram entropy, richly varied letter sequences, while still showing a steeply Zipfian frequency profile, dominated by a handful of generic elements repeated across many otherwise-distinct names, and the companion series' gazetteer pipelines track the two separately for exactly this reason, following the same discipline Section D.2 insisted on for keeping proven and merely empirical findings apart.

D.5 Substrate Layering: A Methodological Limit of the Bigram Window

Section D.2 was careful to mark bigram entropy and inter-zone divergence as empirical findings, rigorous within their statistical framework but not proven in Chapter 4 and Chapter 5's sense; it is worth closing this appendix by naming, explicitly, the specific methodological limit responsible for that empirical status, rather than leaving the caution general. A substrate zone, in the companion series' own definition (Section D.1), is a region grouped by shared linguistic origin; in practice, many real regions carry more than one such layer superimposed, an older substrate of names laid down by an earlier population, overwritten in part but not entirely by a later superstrate's own naming practice, so that a single gazetteer for one modern administrative region can contain toponyms drawn from two or more historically distinct naming traditions, coexisting rather than replacing one another cleanly.

A bigram model of the kind Section D.1 describes, applied to such a mixed gazetteer without first separating its layers, returns a single blended distribution, and the entropy and divergence figures the companion series computes from it describe that blend rather than either underlying tradition on its own — a limitation the companion volumes disclose directly, in keeping with this volume's own discipline of reporting a gap honestly rather than quietly working around it. Layer separation, where attempted at all in the companion series, proceeds by flagging toponyms with etymologies independently attested in the historical-linguistic literature and analyzing the resulting sub-corpora separately, a procedure that depends entirely on the prior availability of that independent etymological attestation and cannot be performed by the bigram method itself, which sees only letter sequences and has no way to distinguish a genuinely layered zone from a homogeneous one built from the same statistics by coincidence. This limit is the toponomastic appendix's own version of Section 9.3's distinction between an imperfect method and an unstable fact: where a substrate zone's layering goes undetected, the resulting entropy figure is not wrong about the corpus actually measured, but it is silent about a structure the corpus contains and the method, unaided, cannot see.

Appendix E — NATRF2022 and the Realization of a Geodetic Datum

E.1 What a Datum Is, and Why a Realization Drifts

Chapter 9 treated the instability of a geodetic datum as a philosophical problem — the ground a coordinate is stated against does not hold still — without dwelling on the specific technical apparatus that makes this so. It is worth supplying that apparatus here, briefly, because the philosophical claim rests on a distinction the technical literature makes precisely: between a datum, an abstract mathematical definition of an origin, an orientation, and a reference ellipsoid, and a realization of that datum, a concrete set of coordinates assigned to physical survey marks at a stated epoch, which is what any working surveyor or GNSS receiver actually consults. The datum itself, defined once, does not drift; its realization does, because the physical marks it is pinned to sit on tectonic plates that move continuously, a few centimeters a year for most of North America, small on a human timescale but cumulative over the decades a coordinate system is expected to remain in service, so that a coordinate correct in one epoch is measurably wrong in another even though nothing about the underlying definition has changed.

The North American Datum of 1983 (NAD 83), the reference frame in use across the United States for four decades, exemplifies the problem Chapter 9 named: it was originally defined with an origin offset from the Earth's true center of mass by roughly two meters, a discrepancy tolerable for terrestrial surveying but increasingly at odds with GNSS positioning, which naturally reports coordinates relative to the geocenter, and its vertical companion, the North American Vertical Datum of 1988 (NAVD 88), depends on a dense network of physical benchmarks that corrode, subside, and are periodically destroyed by the same construction and land use they were meant to support. Neither defect is a failure of the original surveyors' competence; both are the datum-realization gap Section 9.2 already argued is not a technical inconvenience awaiting a better instrument, but

a structural feature of pinning an abstract reference frame to a physical, moving Earth at a particular moment in time.

E.2 NATRF2022: A Plate-Fixed, Epoch-Aware Replacement

The U.S. National Geodetic Survey's modernization of the National Spatial Reference System replaces NAD 83 with four new terrestrial reference frames, of which the North American Terrestrial Reference Frame of 2022 (NATRF2022) is the one covering the bulk of the continent, paired with a new vertical geopotential datum, NAPGD2022, that replaces NAVD 88's benchmark network with a hybrid geoid model tied directly to GNSS heights. The new frames are explicitly plate-fixed: rather than defining a single global set of coordinates and letting plate motion silently erode their accuracy, as NAD 83 effectively did, each of the four replacement frames is defined to co-move with the tectonic plate it covers, and every published coordinate in it carries a stated reference epoch, so that a position is properly given not as a pair of numbers alone but as those numbers together with the date they were true, plus, where needed, a velocity to translate them to another date.

This is a genuine redesign of what "the same point" means administratively, not merely a numerical recalibration of NAD 83's old origin offset. NAD 83 asked a coordinate to be permanently correct; NATRF2022 asks a coordinate to be correct at a stated moment and provides the machinery, an epoch and a velocity field, to make that moment explicit rather than assumed. The National Geodetic Survey's own transition materials warn that this changes published coordinates for the same physical monument by amounts up to several decimeters in places, not because the monument moved further than expected but because the reference frame it is stated against has been redefined to track the motion honestly instead of absorbing it silently into an aging, undated origin — the same distinction between a stable definition and an unstable realization that Section 9.2 drew in the abstract.

E.3 What the Geodetic Case Adds to Chapter 9's Ontological Question

Chapter 9 argued that the instability Section 9.1 identified is not a defect to be engineered away, in exactly the sense Chapter 10 later generalizes for cartographic distortion as a whole, but the NATRF2022 transition sharpens that claim in a way the chapter's more abstract discussion could not: it shows a major national geodetic authority formally institutionalizing the distinction between datum and realization, building the instability directly into the coordinate format itself — requiring an epoch, offering a velocity — rather than treating it as an occasional correction applied when the accumulated drift finally becomes too large to ignore, as the periodic realignments of NAD 83 (2007, 2011) had done before it. Where Chapter 9 asked whether the ground a coordinate is stated against can be assumed to hold still, NATRF2022 answers, in effect, no, and further, that the correct response is not a better fixed frame but an explicitly time-indexed one, a position Chapter 9's own argument anticipates but does not, on its own philosophical terms, need this technical confirmation to establish; the confirmation is offered here as exactly what Chapter 9's Preliminary Note register keeps separate — a fact about current geodetic practice, not a further theorem — supporting, rather than proving, the chapter's ontological point.

E.4 A Comparative Case: Australia's GDA2020 and a Different Institutional Answer to the Same Drift

NATRF2022 is not the only institutional response the drift problem Section E.1 named has received, and a second, independently arrived-at solution sharpens what Section E.3 already argued: that the response is a matter of institutional design, not of geometry. Australia sits on one of the fastest-moving continental plates on Earth, drifting northward at roughly seven centimeters a year, so that by the time its Geocentric Datum of Australia 1994 (GDA94) had been in use for over two decades, the accumulated offset between that plate-fixed datum and the globally realized, satellite-based coordinates ordinary GPS receivers report had grown to nearly two meters, large enough to matter for autonomous machinery, precision agriculture, and other applications that had not existed when GDA94 was defined. Geoscience Australia's replacement, the Geocentric Datum of Australia 2020 (GDA2020), took a structurally different path from NATRF2022's: rather than folding an epoch and a velocity model into the primary legal datum itself, GDA2020 remains, like its predecessor, a static, plate-fixed frame, and the time-dependence Chapter 9 identified as unavoidable is instead pushed into a separate, explicitly dynamic reference frame, the Australian Terrestrial Reference Frame 2014 (ATRF2014), and a published transformation grid, consulted only by the applications that need agreement with real-time satellite coordinates.

Two national geodetic authorities, facing what Chapter 9 established is the same forced fact — a plate-fixed datum and a satellite-realized position will not stay reconciled indefinitely — arrived at materially different architectures for accommodating it: the United States built the instability into NATRF2022's primary coordinate format, requiring an epoch of every user; Australia kept GDA2020's primary format static and unchanged, and confined the instability to a secondary layer most users never touch. Nothing in Chapter 9's argument, or in Section E.1's restatement of it, favors one architecture over the other; both are compliant accommodations of the same ontological fact, and the choice between folding time-dependence into the primary datum or into a secondary transformation is exactly the kind of residual, non-forced freedom Section 6.3

already described for geometric convention in general, now instantiated a second time in the specific vocabulary of geodesy rather than cartography.

Appendix F — Mercator, Peters, and the Cartography of Power

F.1 Mercator's 1569 Projection: A Navigational Tool, Not a Political One

Chapter 13 argued for a distinction Section 13.1 stated as the chapter's central discipline: between distortion that is mathematically forced, in Chapter 4's proven sense, and distortion that is politically chosen, a further, non-geometric fact about who selected a given projection and why. Gerardus Mercator's 1569 world map is the case around which the modern debate over that distinction has organized itself, and it is worth setting out its own history plainly before Section F.2 turns to its most famous critic. Mercator designed his projection for a single, stated, and narrowly practical purpose: to let a navigator plot a constant compass bearing as a straight line on the chart, a property, conformality, Chapter 4's own theorem already proves cannot be had except at the cost of area, and which becomes severe at high latitudes, where Mercator's own construction inflates Greenland to a size visually comparable to Africa, though Africa's true area is roughly fourteen times greater.

The area distortion, in other words, is not incidental to Mercator's purpose; it is the specific, provable price Chapter 4's theorem exacts for the specific property Mercator needed, and needed for a reason that had nothing to do with the relative visual prominence of any continent. This is exactly the forced half of Section 13.1's distinction: no alternative construction preserving compass bearings avoids it, and choosing conformality for celestial and marine navigation, as sixteenth-century Europe's expanding maritime traffic demanded, is not itself a political act in the sense Chapter 13 goes on to examine, whatever political uses the resulting map was later put to in classrooms and boardrooms centuries after its stated purpose had receded from view.

F.2 The Peters Projection and the Equal-Area Counter-Claim

In 1973, the German historian Arno Peters presented a competing world map built on an equal-area projection — a construction, in fact, substantially identical to one James Gall had published in 1855, which is why cartographers now generally call it the Gall-Peters projection — explicitly as a corrective to what Peters argued was the Mercator map's Eurocentric distortion of relative size, and promoted it through explicitly political channels, adopted rapidly by development and relief organizations on the strength of its claim to show the Global South at its true proportional area. The equal-area property itself is, again, exactly Chapter 4's theorem in action: Tissot's indicatrix, Section 4.2 already showed, cannot be a uniform circle everywhere on a flat map, and a projection that fixes area exactly, as Gall-Peters does, pays for it in the opposite currency Mercator paid in, severe shearing of shape rather than of area, visible in the elongated, compressed silhouettes of continents near the equator and poles alike that a critic in Chapter 13's own source material compared to "long winter underwear hung out to dry on the Arctic Circle."

What makes the Gall-Peters case sharper than Mercator's for Section 13.1's purposes is that Peters did not merely choose a projection for a stated practical end, as Mercator had; he marketed his choice explicitly as more accurate than its rivals, a claim Chapter 4's theorem makes precise and rules out for any flat projection whatsoever, equal-area included. The professional cartographic community's rejection of this specific claim — formalized in a 1989 resolution by seven North American geographic organizations denouncing the promotion of any single rectangular world map as uniquely correct — was not, on this volume's own reading, a rejection of Peters' underlying political argument that the Mercator map had been used, and often still is used, in contexts where its navigational conformality serves no purpose and its polar-region inflation quietly shapes a viewer's sense of relative scale; it was a rejection of the further, specifically false claim that any projection escapes Chapter 4's trade-off altogether.

F.3 What Was Mathematically Forced and What Was Rhetorically Chosen

Laid side by side, the two cases let Section 13.1's distinction be stated with unusual precision. What was forced, in both maps, is exactly what Chapter 4 proves forced for any flat projection of a sphere: some distortion, distributed according to whichever property, angle or area, the cartographer chose to preserve exactly. What was chosen, and chosen twice over in each case, is first the property itself — Mercator's conformality, serving a stated navigational purpose; Peters' equal area, serving a stated corrective, political purpose — and second, and this is the choice Chapter 13's critical-cartography discussion concerns itself with, the further decisions about where each map was subsequently deployed, in whose classrooms, on whose walls, for audiences with no navigational or areal-comparison need being served at all. Harley's argument, engaged in Chapter 13's main text, that a map is never merely a neutral geometric object but always also a rhetorical one, holds for both cases without needing either map's forced distortion to itself be the political fact; the political fact, in both, is downstream of the mathematics, not identical with it.

This appendix's closing observation, and the reason the Mercator/Peters pair belongs here rather than folded into Chapter 13's main argument, is that the pair shows Section 13.1's distinction surviving contact with the single most publicly contested case in the entire history of cartography, a case where a critic's political argument (Peters') and a mathematician's provable theorem (Chapter 4's) point in genuinely different directions and are, on inspection, both right about what they each actually claim: Peters was right that a specific historical deployment of Mercator's map carried an unexamined, Eurocentric visual rhetoric; the cartographic establishment was right that no projection, Peters' own included, escapes the theorem Chapter 4 proves; and neither correctness cancels the other, because, as Section 13.1 insisted from the outset, they are answers to two different questions wearing the same map's clothing.

F.4 A Twenty-First-Century Coda: Web Mercator and the Same Debate in a New Medium

The Mercator/Peters controversy Sections F.1 through F.3 traced in print did not end with print; it resurfaced largely unchanged when online mapping platforms, Google Maps foremost among them, adopted a spherical variant of Mercator's construction, known as Web Mercator (EPSG:3857), as the near-universal standard for zoomable web maps beginning in the mid-2000s. The choice was made for reasons wholly internal to the engineering problem of serving map tiles over the internet rather than for any navigational purpose Mercator himself would have recognized: conformality lets small square tiles be rendered, cached, and stitched together at any zoom level without local shape distortion inside a given tile, and Mercator's projection happens to produce exactly the fixed-aspect-ratio tiles the standard 256-by-256-pixel raster scheme underlying most web map software requires. The visible cost, at the fully zoomed-out, whole-world scale, reproduces precisely the areal exaggeration Peters had objected to in print, Greenland and northern Europe rendered far larger relative to equatorial regions than their true areas warrant, and the criticism resurfaced accordingly, directed this time at Google specifically, until the company announced on August 2, 2018 that Google Maps, on desktop only, would switch, at the single most zoomed-out scale only, to a three-dimensional globe view rather than the flat Web Mercator projection, while retaining Web Mercator at every more zoomed-in scale, and on mobile entirely, where the tiling requirement still governs.

Read through Section 13.1's distinction, which this appendix's main text has already applied once to Mercator and Peters, Google's 2018 change is a clean illustration rather than a complication of the same point made in a new medium: nothing in the tiling requirement that genuinely justifies Web Mercator forces a flat projection specifically at the single, whole-world zoom level, and Chapter 4's theorem forces only that some distortion exist somewhere, never that it be tolerated at every scale a product happens to offer. The pattern this appendix has now traced across four centuries and two media, sixteenth-century print and

twenty-first-century browser, is the same throughout: the mathematics forces a floor beneath which distortion cannot fall, and everything built above that floor, including four decades of accepting a particular exaggeration at global scale and then, in 2018, declining to any longer, remains a historically contingent and institutionally revisable choice, not a further theorem.

Appendix G — Arrow's Theorem as a Second Combinatorial Impossibility

G.1 Arrow's Impossibility Theorem, Stated

Section 5.8 invoked Kenneth Arrow's impossibility theorem in passing, as an example, alongside Chapter 5's own hairy ball obstruction, of a proof that a global requirement cannot be met without a forced, quantifiable exception; it is worth stating the theorem itself here, since the main chapters never do, so that the analogy Section 5.8 draws can be checked rather than merely asserted. Arrow's theorem, first published in his 1951 monograph Social Choice and Individual Values, concerns the aggregation of individual preference orderings over three or more options into a single social ordering. It shows that no such aggregation procedure can simultaneously satisfy a short list of conditions that, taken individually, each look like a minimal requirement of fairness: that the procedure accept any possible combination of individual preferences (unrestricted domain), that a unanimous preference for one option over another always be reflected in the social ordering (weak Pareto), that the social ranking of any two options depend only on individual rankings of those two options and not on a third (independence of irrelevant alternatives), and that no single individual's preferences alone determine the social ordering regardless of everyone else's (non-dictatorship). Arrow proved that for any domain of three or more options, these conditions are jointly unsatisfiable: some individual must be a dictator, or one of the other conditions must fail.

The proof is combinatorial rather than topological or geometric: it proceeds by constructing, from the assumption that all four conditions hold, a chain of forced logical consequences that ends by identifying a single individual whose strict preference between any pair always determines the social outcome, regardless of every other individual's preferences — precisely the dictator the fourth condition was meant to rule out. Nothing in the argument invokes continuity, curvature, or any structure from analysis; the impossibility is a fact about finite combinatorics of orderings, which is exactly what makes it, as Section 5.8 suggested, a genuinely independent instance of the pattern Chapter 5 proves geometrically rather than a disguised restatement of it.

G.2 A Structural Parallel to Chapter 5, Not an Identity

The parallel Section 5.8 draws, and the reason Arrow's theorem belongs in a volume about cartographic impossibility at all, is structural in the same sense Appendix A.4 already used for the twelve-pentagon result: both Chapter 5's hairy ball obstruction and Arrow's theorem are proofs that a locally reasonable requirement — a continuous choice of tangent direction at every point of a sphere; a fair aggregation of preferences for every possible input — cannot be extended to a globally consistent one without a forced exception, a vanishing point in one case, a dictator in the other. In each case the forced exception is not a sign of a poorly chosen construction, correctable by a cleverer field or a smarter voting rule; it is proven, by an argument that examines every candidate construction at once, to be unavoidable given the stated conditions. And in each case, exactly as Chapter 6's conventionalism argues for the cartographic case, the theorem itself is silent about which specific exception to accept — where the hairy ball's vanishing point should sit; which of Arrow's four conditions a working voting system should be willing to relax — leaving a residual, non-geometric choice that real electoral systems make in practice by weakening independence of irrelevant alternatives far more often than the other three.

This is the same two-part shape Chapter 11 found, empirically rather than by proof, across three philosophical accounts of scientific representation, and the

same shape Section 5.8 flagged without elaborating: a forced impossibility, stated and proven with full rigor, paired with a genuinely free choice among the ways of living with it. Arrow's case strengthens the claim that this two-part shape is not peculiar to surfaces and curvature, since nothing about vector fields, spheres, or Euler characteristics appears anywhere in Arrow's proof; it is a second, independently discovered instance of the same abstract pattern, arrived at by economists studying voting rather than by geometers studying maps.

G.3 Where the Parallel Breaks Down

This volume's own discipline, stated in the Preliminary Note and enforced throughout, requires marking exactly where such a parallel stops holding, and Arrow's case has a disanalogy worth stating as plainly as the resemblance. Chapter 5's theorem is a statement about continuous vector fields on a specific topological object, the two-sphere, and its proof, whether Milnor's analytic version or the original Poincaré–Hopf argument, is inseparable from that object's Euler characteristic; change the surface to a torus and the obstruction vanishes entirely, a continuous nonvanishing tangent field exists there without difficulty. Arrow's theorem carries no such surface-specific escape clause: it holds for any domain of three or more options whatsoever, with no analogue of "move to a torus" available, because nothing in its combinatorial proof depends on a geometric structure that a different choice of underlying space could alter. The two theorems are, in this precise sense, impossibility results of different logical character, one topological and object-specific, the other combinatorial and domain-general, and Section 5.8's analogy should be read, in keeping with this volume's own register discipline, as pointing to a shared two-part shape — forced impossibility plus residual choice — rather than as claiming any deeper mathematical kinship between the hairy ball obstruction and Arrow's dictator.

G.4 A Practical Postscript: How Voting Systems Live With Arrow's Theorem Anyway

Section 4.7's discussion of compromise projections has a close structural cousin in the practical history of voting-system design, worth stating briefly as this appendix's closing case. Arrow's theorem, as Section G.1 stated it, proves that no ranked voting rule can simultaneously satisfy all four of Arrow's named fairness conditions for three or more candidates; every voting system actually in use, plurality, instant-runoff (ranked-choice) voting, the Borda count, approval voting, and the various Condorcet methods among them, responds not by evading the theorem, which Sections G.2 and G.3 have already established cannot be done, but by openly sacrificing exactly one of the four conditions, a choice electoral reformers make explicitly rather than by accident. Instant-runoff voting is well documented to violate independence of irrelevant alternatives, since adding or removing a losing candidate can, in certain configurations, change which of the remaining candidates wins, a cost its advocates accept in exchange for eliminating the vote-splitting, or 'spoiler,' effect that plain plurality voting suffers from instead.

Read against Section 4.3's account of choosing among Tissot's survivors, the parallel is close to exact. Just as no projection can be simultaneously conformal and equal-area, so no ranked voting rule can satisfy every one of Arrow's conditions, and in both cases the practical response has been neither paralysis nor denial but a menu of named, well-understood options, each openly advertising which property it has given up, Mercator giving up equal area for conformality, instant-runoff voting giving up independence of irrelevant alternatives for protection against vote-splitting, with the choice among them left to the institution adopting the system rather than settled by the proof. Arrow's theorem no more forces plurality voting on a democracy than Chapter 4's theorem forces the Mercator projection on a cartographer; both leave a genuine, consequential choice standing exactly where the proof says a choice must be made, and this appendix closes, appropriately, by finding in electoral design the

same two-part shape, forced impossibility paired with real residual freedom, that this volume's cartographic case has traced from its first chapter to its last.

Bibliography

The chapters and appendices of this volume cite, by name, a number of primary results and historical sources established or reported elsewhere — in the companion technical monographs this volume reinterprets philosophically, and in the wider literature on cartography, topology, geodesy, and the philosophy of science. The list below gathers the latter: works referred to by name in the main text or the appendices that are not themselves companion volumes of this project. It is organized alphabetically by author and makes no claim to completeness beyond what this volume actually cites; it is offered as a reader's convenience, not as a survey of any of the fields it touches.

Adams, Oscar S. (1925). "A Study of Map Projections in General." U.S. Coast and Geodetic Survey Special Publication No. 88, presenting the hemisphere-in-a-square projection.

Akaike, Hirotugu (1974). "A New Look at the Statistical Model Identification." IEEE Transactions on Automatic Control, 19(6), 716–723.

Arrow, Kenneth J. (1951). Social Choice and Individual Values. New York: John Wiley & Sons. (2nd ed. 1963, New Haven: Yale University Press.)

Bateson, Gregory (1972). Steps to an Ecology of Mind. San Francisco: Chandler Publishing.

Batterman, Robert W. (2002). "Asymptotics and the Role of Minimal Models." The British Journal for the Philosophy of Science, 53(1), 21–38.

Baudrillard, Jean (1981). Simulacres et Simulation. Paris: Éditions Galilée. (Trans. Sheila Faria Glaser as Simulacra and Simulation, 1994, Ann Arbor: University of Michigan Press.)

Bergson, Henri (1889). Essai sur les données immédiates de la conscience. Paris: Félix Alcan. (Trans. F.L. Pogson as Time and Free Will: An Essay on the Immediate Data of Consciousness, 1910, London: George Allen & Unwin.)

Borges, Jorge Luis (1946). “Del rigor en la ciencia.” Los Anales de Buenos Aires, 1(3). (Trans. as “On Exactitude in Science,” in Collected Fictions, 1998, New York: Viking.)

Borsuk, Karol (1933). “Drei Sätze über die n-dimensionale euklidische Sphäre.” Fundamenta Mathematicae, 20, 177–190. (States the theorem now known as the Borsuk–Ulam theorem, crediting a conjecture of Stanisław Ulam.)

Carnap, Rudolf (1934). Logische Syntax der Sprache. Vienna: Julius Springer. (Trans. Amethe Smeaton as The Logical Syntax of Language, 1937, London: Kegan Paul, Trench, Trubner.)

Cartwright, Nancy (1999). The Dappled World: A Study of the Boundaries of Science. Cambridge: Cambridge University Press.

Chebyshev, Pafnuty L. (1856). “Sur la construction des cartes géographiques.” Bulletin de la Classe Physico-Mathématique de l'Académie Impériale des Sciences de Saint-Pétersbourg, 14, 257–261.

Coxeter, H.S.M. (1979). “The Non-Euclidean Symmetry of Escher's Picture ‘Circle Limit III.’” Leonardo, 12, 19–25, 32.

Craig, James Ireland (1910). “Map Projections for Meteorological Charts.” Survey of Egypt, presenting the retroazimuthal (“Mecca”) projection of 1909.

Deleuze, Gilles, and Guattari, Félix (1980). Mille Plateaux: Capitalisme et schizophrénie 2. Paris: Éditions de Minuit. (Trans. Brian Massumi as A Thousand Plateaus: Capitalism and Schizophrenia, 1987, Minneapolis: University of Minnesota Press.)

Dent, Steve (2018). “Google Maps Now Zooms Out to a Globe Instead of a Flat Earth.” Engadget, August 3, 2018 (reporting Google Maps’ desktop-only switch to a three-dimensional globe view at the fully zoomed-out scale).

Duhem, Pierre (1906). La Théorie physique: son objet et sa structure. Paris: Chevalier & Rivière. (Trans. Philip P. Wiener as The Aim and Structure of Physical Theory, 1954, Princeton: Princeton University Press.)

Eco, Umberto (1994). “On the Impossibility of Drawing a Map of the Empire on a Scale of 1 to 1.” In How to Travel with a Salmon and Other Essays, trans. William Weaver. New York: Harcourt Brace.

Elgin, Catherine Z. (2017). True Enough. Cambridge, MA: MIT Press.

Eliade, Mircea (1957). Das Heilige und das Profane. Reinbek bei Hamburg: Rowohlt. (Trans. Willard R. Trask as The Sacred and the Profane: The Nature of Religion, 1959, New York: Harcourt, Brace.)

Eyring, Veronika, Bony, Sandrine, Meehl, Gerald A., Senior, Catherine A., Stevens, Bjorn, Stouffer, Ronald J., and Taylor, Karl E. (2016). "Overview of the Coupled Model Intercomparison Project Phase 6 (CMIP6) Experimental Design and Organization." Geoscientific Model Development, 9(5), 1937–1958.

Gall, James (1855). "On Improved Monographic Projections of the World." Presented to the British Association for the Advancement of Science, Glasgow (equal-area projection later revived by Peters).

Gauss, Carl Friedrich (1827). "Disquisitiones generales circa superficies curvas." Commentationes Societatis Regiae Scientiarum Gottingensis, 6, 99–146.

Giere, Ronald N. (2006). Scientific Perspectivism. Chicago: University of Chicago Press.

Goodman, Nelson (1968). Languages of Art: An Approach to a Theory of Symbols. Indianapolis: Bobbs-Merrill.

Grave, Dmitry A. (1896). Ob osnovnykh zadachakh matematicheskoy teorii postroeniya geograficheskikh kart [On the Basic Problems of the Mathematical Theory of the Construction of Geographic Maps]. Doctoral dissertation, St. Petersburg, giving the first published proof of Chebyshev's 1856 theorem.

Guyou, Émile (1887). "Nouveau système de projection de la sphère: Généralisation de la projection de Mercator." Annales Hydrographiques, Sér. 2, 9, 16–35.

Hammer, Ernst (1910). "Unechte azimutale Entfernungsprojektion." Petermanns Geographische Mitteilungen, 56, presenting the Hammer retroazimuthal projection.

Harley, J.B. (1989). "Deconstructing the Map." Cartographica: The International Journal for Geographic Information and Geovisualization, 26(2), 1–20.

Harrison, Craig, Brown, Nicholas, Dawson, John, and Fraser, Roger (2023). "Geocentric Datum of Australia 2020: The First Australian Datum Developed from a Rigorous Continental-Scale Adjustment." Journal of Spatial Science. DOI: 10.1080/14498596.2023.2184429.

Heidegger, Martin (1927). Sein und Zeit. Halle: Max Niemeyer. (Trans. John Macquarrie and Edward Robinson as Being and Time, 1962, Oxford: Blackwell.)

Hesse, Mary (1966). Models and Analogies in Science. Notre Dame: University of Notre Dame Press.

Hopf, Heinz (1927). "Vektorfelder in n-dimensionalen Mannigfaltigkeiten." Mathematische Annalen, 96(1), 225–249.

Husserl, Edmund (1934/1981). "Foundational Investigations of the Phenomenological Origin of the Spatiality of Nature," trans. Fred Kersten. In Husserl: Shorter Works, ed. Peter McCormick and Frederick A. Elliston, 222–233. Notre Dame: University of Notre Dame Press. (Manuscript usually known by its working title, 'The Earth Does Not Move,' written in 1934 and unpublished in Husserl's lifetime.)

Kant, Immanuel (1781/1787). Kritik der reinen Vernunft. Riga: Johann Friedrich Hartknoch. (Trans. Norman Kemp Smith as Critique of Pure Reason, 1929, London: Macmillan.)

Kolmogorov, Andrei N. (1965). "Three Approaches to the Quantitative Definition of Information." Problems of Information Transmission, 1(1), 1–7.

Korzybski, Alfred (1933). Science and Sanity: An Introduction to Non-Aristotelian Systems and General Semantics. Lakeville, CT: Institute of General Semantics.

Kuhn, Thomas S. (1962). The Structure of Scientific Revolutions. Chicago: University of Chicago Press.

Latour, Bruno (1986). "Visualization and Cognition: Thinking with Eyes and Hands." Knowledge and Society: Studies in the Sociology of Culture Past and Present, 6, 1–40.

Lewis, David (1969). Convention: A Philosophical Study. Cambridge, MA: Harvard University Press.

Littrow, Joseph Johann von (1833). Untitled construction of the conformal retroazimuthal projection, Vienna. (Independently rediscovered by Patrick Weir, 1890.)

Lynch, Kevin (1960). The Image of the City. Cambridge, MA: MIT Press.

Marr, David (1982). Vision: A Computational Investigation into the Human Representation and Processing of Visual Information. San Francisco: W.H. Freeman.

Merleau-Ponty, Maurice (1945). Phénoménologie de la perception. Paris: Gallimard. (Trans. Colin Smith as Phenomenology of Perception, 1962, London: Routledge & Kegan Paul.)

Milnor, John W. (1978). “Analytic Proofs of the ‘Hairy Ball Theorem’ and the Brouwer Fixed Point Theorem.” The American Mathematical Monthly, 85(7), 521–524.

National Geodetic Survey (2022–2026). “New Datums: NATRF2022 and NAPGD2022.” NOAA National Ocean Service. Federal Register notices on the modernization of the National Spatial Reference System.

National Public Radio (2017). “Boston Students Get a Glimpse of a Whole New World, With Different Maps.” The Two-Way, March 21, 2017 (reporting Boston Public Schools’ March 2017 adoption of the Gall–Peters projection for classroom maps).

Pareto, Vilfredo (1906). Manuale di economia politica. Milan: Società Editrice Libraria.

Peirce, Charles S. (1879). “A Quincuncial Projection of the Sphere.” American Journal of Mathematics, 2(4), 394–397.

Peters, Arno (1974). “Die neue Kartographie / The New Cartography.” Lecture before the Deutsche Gesellschaft für Kartographie, presenting the 1973 projection subsequently known as Gall-Peters.

Poincaré, Henri (1902). La Science et l'Hypothèse. Paris: Flammarion.

Potochnik, Angela (2017). Idealization and the Aims of Science. Chicago: University of Chicago Press.

Ptolemy, Claudius (c. 150 CE). Geōgraphikē Hyphēgēsis. Trans. J. Lennart Berggren and Alexander Jones as Ptolemy's Geography: An Annotated Translation of the Theoretical Chapters, 2000, Princeton: Princeton University Press.

Quine, Willard V.O. (1951). “Two Dogmas of Empiricism.” The Philosophical Review, 60(1), 20–43.

Reichenbach, Hans (1928). Philosophie der Raum-Zeit-Lehre. Berlin: Walter de Gruyter. (Trans. Maria Reichenbach as The Philosophy of Space and Time, 1958, New York: Dover.)

Robinson, Arthur H. (1974). “A New Map Projection: Its Development and Characteristics.” International Yearbook of Cartography, 14, 145–155.

Šavrič, Bojan, Patterson, Tom, and Jenny, Bernhard (2018). “The Equal Earth Map Projection.” International Journal of Geographical Information Science, 33(3), 454–465.

Simon, Herbert A. (1947). Administrative Behavior: A Study of Decision-Making Processes in Administrative Organization. New York: Macmillan.

Snyder, John P. (1993). Flattening the Earth: Two Thousand Years of Map Projections. Chicago: University of Chicago Press.

Snyder, John P., and Voxland, Philip M. (1989). An Album of Map Projections. U.S. Geological Survey Professional Paper 1453. Washington, D.C.: United States Government Printing Office.

Strevens, Michael (2008). Depth: An Account of Scientific Explanation. Cambridge, MA: Harvard University Press.

Suárez, Mauricio (2004). "An Inferential Conception of Scientific Representation." Philosophy of Science, 71(5), 767–779.

Tissot, Nicolas-Auguste (1881). Mémoire sur la représentation des surfaces et les projections des cartes géographiques. Paris: Gauthier-Villars.

Tuan, Yi-Fu (1974). Topophilia: A Study of Environmental Perception, Attitudes, and Values. Englewood Cliffs, NJ: Prentice-Hall.

van Fraassen, Bas C. (1980). The Scientific Image. Oxford: Clarendon Press.

Véronneau, Marc, and Craymer, Michael (2019). "The U.S. is Replacing NAD83 with NATRF2022: What This Means for Canada." Geomatica, Canadian Geodetic Survey.

Wimsatt, William C. (2007). Re-Engineering Philosophy for Limited Beings: Piecewise Approximations to Reality. Cambridge, MA: Harvard University Press.

Wittgenstein, Ludwig (1922). Tractatus Logico-Philosophicus, trans. C.K. Ogden. London: Kegan Paul, Trench, Trubner.

Zipf, George Kingsley (1949). Human Behavior and the Principle of Least Effort: An Introduction to Human Ecology. Cambridge, MA: Addison-Wesley Press.